Ecological Applications, 26(8), 2016, pp. 2677–2694

Fitting state-space integral projection models to size-structured time series data to estimate unknown parameters

J. Wilson White,[1,7] Kerry J. Nickols,[2] Daniel Malone,[3] Mark H. Carr,[3] Richard M. Starr,[4] Flora Cordoleani,[5] Marissa L. Baskett,[6] Alan Hastings,[6] and Louis W. Botsford[5]

[1]*Department of Biology and Marine Biology, University of North Carolina Wilmington, Wilmington, North Carolina 28043 USA*
[2]*Division of Science and Environmental Policy, California State University Monterey Bay, Seaside, California 93955 USA*
[3]*Department of Ecology and Evolutionary Biology, University of California Santa Cruz, Santa Cruz, California 95060 USA*
[4]*California Sea Grant Extension Program, Moss Landing Marine Laboratories, Moss Landing, California 95039 USA*
[5]*Department of Wildlife, Fish, and Conservation Biology, University of California Davis, Davis, California 95616 USA*
[6]*Department of Environmental Science and Policy, University of California Davis, Davis, California 95616 USA*

Abstract. Integral projection models (IPMs) have a number of advantages over matrix-model approaches for analyzing size-structured population dynamics, because the latter require parameter estimates for each age or stage transition. However, IPMs still require appropriate data. Typically they are parameterized using individual-scale relationships between body size and demographic rates, but these are not always available. We present an alternative approach for estimating demographic parameters from time series of size-structured survey data using a Bayesian state-space IPM (SSIPM). By fitting an IPM in a state-space framework, we estimate unknown parameters and explicitly account for process and measurement error in a dataset to estimate the underlying process model dynamics. We tested our method by fitting SSIPMs to simulated data; the model fit the simulated size distributions well and estimated unknown demographic parameters accurately. We then illustrated our method using nine years of annual surveys of the density and size distribution of two fish species (blue rockfish, *Sebastes mystinus*, and gopher rockfish, *S. carnatus*) at seven kelp forest sites in California. The SSIPM produced reasonable fits to the data, and estimated fishing rates for both species that were higher than our Bayesian prior estimates based on coast-wide stock assessment estimates of harvest. That improvement reinforces the value of being able to estimate demographic parameters from local-scale monitoring data. We highlight a number of key decision points in SSIPM development (e.g., open vs. closed demography, number of particles in the state-space filter) so that users can apply the method to their own datasets.

Key words: fishing rate; integral projection model; particle filter; Sebastes carnatus; Sebastes mystinus; *state-space model.*

Introduction

Structured population models have a long history in conservation and natural resources management, from population viability analysis to fisheries stock assessment, because they allow understanding of how age-, stage-, or size-specific anthropogenic impacts affect population dynamics and management outcomes (Crouse et al. 1987, Doak et al. 1994, Methot and Wetzel 2013). However, analysis of structured population dynamics using matrix methods (e.g., Caswell 2001) has large data requirements and may depend on unrealistic assumptions, so integral projection models (IPMs; Easterling et al. 2000) are an increasingly popular tool for analyzing population dynamics to test hypotheses regarding persistence, geographical distributions, and other emergent properties (Coulson 2012, Merow et al. 2014). An IPM is conceptually similar to a traditional age- or size-structured matrix population model (such as those described by Caswell 2001), except that the IPM describes the population in terms of a continuous distribution over size rather than abundance within discrete size bins. Thus, the IPM uses a smooth kernel (rather than a discrete projection matrix) to describe the probability of transitioning between sizes. The kernel is comprised of continuous functions describing size-specific growth, fecundity, and mortality rates. Consequently, instead of requiring parameter estimates of every possible stage transition probability, one need only estimate the parameters of those continuous functions, thus achieving a substantial savings in parameter estimation. Once the kernel has been estimated, the behavior of the IPM can be analyzed in much the same way as traditional matrix projection models, such as examining the dominant eigenvalue and eigenvector to determine the long-term asymptotic growth rate and population distribution, and calculating the sensitivities of the eigenvalue to the various demographic parameters (Easterling et al. 2000, Ellner and Rees 2006*a, b*, Rees and Ellner 2009). To date, most published implementations of IPMs have been used in this fashion, either to examine asymptotic long-term growth

Manuscript received 27 October 2015; revised 23 April 2016; accepted 22 June 2016. Corresponding Editor: C. Merow.
[7]E-mail: whitejw@uncw.edu

rates (e.g., Zuidema et al. 2010, Coulson 2012, Madin et al. 2012) or to simulate steady-state dynamics (e.g., Bruno et al. 2011).

Even with the inherent advantages of the IPM approach, the challenge of going from data to model is still substantial. IPM construction typically involves assembling information from organism-scale relationships between size and demographic rates to build the IPM, then projecting the IPM to obtain population-scale size distributions (Merow et al. 2014). There are two potential challenges for this approach. First, individual-scale data are not always available for relevant demographic rates, and some parameters might not be known at all. For example, in marine fisheries it can be difficult to obtain estimates of the fishery harvest rate, particularly at small spatial scales relevant to place-based management (Eero et al. 2012, Castrejón and Charles 2013), and for long-lived tree species, it can be difficult to estimate survival reliably (Ghosh et al. 2012). Second, the kernel based on individual-scale parameters may not accurately capture population-scale processes. This mismatch can arise because individual-scale processes do not smoothly aggregate to represent population-scale dynamics, particularly if demographic rates are nonlinear and sensitive to spatial variation in density, or if there is strong individual heterogeneity (Clark 2003, Clark et al. 2011, Chesson 2012, Ghosh et al. 2012).

A possible remedy to these challenges is to fit the IPM directly to a time series of size distributions in order to estimate unknown parameters for the kernel. This approach allows IPM users to take advantage of the many long-term, size-structured census datasets that are available, particularly for fisheries and terrestrial plants. Such data lend themselves to the size-structured IPM framework, and estimating population parameters in this way can avoid the scale mismatch problem (Ghosh et al. 2012). For example, González and Martorell (2013) obtained maximum-likelihood estimates of kernel parameters for an IPM of a long-lived cactus by fitting model simulations to a 12-year time series of field observations (see also González et al. 2016 for a similar approach). Ghosh and colleagues (Ghosh et al. 2012, Gelfand et al. 2013) developed a different approach to fit an IPM to size-structured survey data (they used a forest tree dataset as an example), and their hierarchical Bayesian method had the key distinction of being a state-space model. That is, Ghosh et al. (2012) assumed the IPM described the underlying true deterministic dynamics of the system, but that the population was also affected by stochastic process error, and the observed data also included measurement error. Using a state-space approach to partition the latent process model dynamics from both process and measurement error is key to obtaining parameter estimates that are useful for future projections (de Valpine and Hastings 2002, Dennis et al. 2006). Fitting a state-space IPM to survey data has great promise for ecological inference but can entail serious computational challenges. Indeed, Ghosh et al. (2012) implemented several clever approximations in order to avoid the computationally intensive direct calculation of Bayesian posterior densities for the parameters in their state-space IPM.

We describe a new approach for constructing state-space IPMs (SSIPMs): We utilize a particle filter (Gordon et al. 1993, Knape and de Valpine 2012) to link the process model to observations and a Bayesian Markov chain Monte Carlo (MCMC) algorithm to estimate unknown parameters. Our approach permits the estimation of uncertain local demographic parameters from a time series of population observations, while accounting for both process and measurement error in those data, and overcomes some limitations faced by prior efforts (Ghosh et al. 2012, González and Martorell 2013, González et al. 2016). We describe our methods, demonstrate the accuracy of our approach using simulated data, and provide an example application using data for two kelp forest fish species surveyed inside and outside of a California marine protected area (MPA). To enhance our readers' ability to adapt this method to their own datasets and applications, throughout this study we identify key decision points in both model construction and model fitting, describe the options available at each decision point, and explain the choices we made for our own implementation. In Fig. 1, we illustrate our overall state-space approach and where those decision points fit in.

Methods

General state-space integral projection model

The basic premise of a state-space model (de Valpine and Hastings 2002, Dennis et al. 2006) is that there is an underlying process model representing the true dynamics of the system, $N(t)$, but that true state is hidden and only observed imperfectly as data, $X(t)$, ($t = 1, 2, \ldots T$). The hidden state evolves over time as a Markov process, with the value at time $t + 1$, $N(t + 1)$, depending only on the state at time t and a process error term, $\nu(t)$: $N(t + 1) = G(N[t], \nu[t])$. Similarly, the observed data at time t depend on the hidden (actual) state at that time and a measurement error term, $\varepsilon(t)$: $X(t) = H(N[t], \varepsilon[t])$. A state-space model uses a filter to estimate the true hidden states $N(t)$ given the observations $X(t)$, and possibly also to estimate unknown parameters of the functions G and H (see Table 1 for list of symbols used in this study).

In our case, N is the true abundance and size structure of individuals at t and X is the corresponding survey observations of those individuals. The function H represents the mechanics of the observation process; for example, we might need to exclude size classes that are not detected in the survey (see *Decision point: observation ogive*). The process model, the function G, is an IPM.

An IPM tracks the state of a population in terms of its size distribution, $N(y,t)$, which is the density function of individuals of size y at time t. Because this is a continuous distribution over y, the actual abundance of individuals

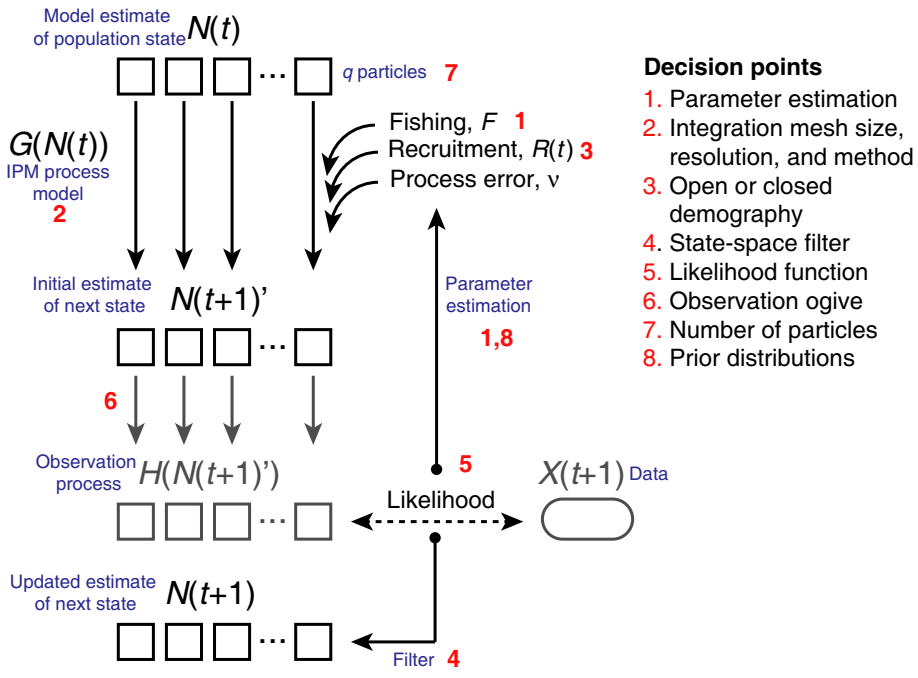

FIG. 1. Schematic of the state-space integral projection model (SSIPM) approach and decision points. One step of the model, the transition from time t to time $t + 1$ is shown. The components of the process model (the IPM) are shown in black text and black boxes; each box representing one particle (i.e., one possible state of the population at that time). Gray text and boxes represent the simulated observations based on the process model, $H(N)$, as well as the actual data, X. Blue text labels the steps of the state-space procedure, and red numbers indicate decision points in the order they appear in the text.

of size y at t (or population density, i.e., number of individuals of size y per unit area) is calculated as the abundance within a small size interval Δy centered on size y: $N(y,t)\Delta y$. The total population size $N(t)$ is $\int_\Omega N(y,t)dy$, where the integral is taken over all biologically reasonable sizes Ω.

In a typical, non-state-space IPM, the population density at time $t + 1$ is the current population density multiplied by the probability density of moving from size x to size y, $K(y,x)$ and integrated over Ω (Easterling et al. 2000):

$$N(y,t+1)=\int_\Omega K(y,x)N(x,t)dx \qquad (1)$$

$K(y,x)$ is known as the kernel and is analogous to the projection matrix in typical structured population models (Caswell 2001). The kernel differs from a projection matrix in that instead of having discrete ages, stages, or size bins with distinct transition probabilities between each pair, the kernel is defined as a combination of functions that are continuous over size, x. Therefore, the kernel can describe the transition between any size and any other size without the coarse binning used in structured matrix models.

To create a generic state-space IPM (SSIPM), we modify Eq. 1 to include a process error term, $\nu(y,t)$, that represents deviations from the projected density due to variation in survival, growth, or reproduction

$$N(y,t+1)=\int_\Omega K(y,x)N(x,t)dx+\nu(y,t). \qquad (2)$$

Eq. 2 serves as function G in the generic state-space model described previously.

We now explain the specific SSIPM implementation we developed. We begin by describing the data in the California MPA case study which provides an illustration of how to apply our method, because the specifics of that system shape some modeling decisions. We then describe the construction of the kernel and other components of the SSIPM, and finally describe the technical aspects of the SSIPM fitting and parameter estimation. In addition to the California MPA dataset, we also generated simulated datasets with known demographic parameters and dynamics in order to validate the SSIPM parameter estimation.

Case study: a California marine protected area

A SSIPM can be fit to any dataset consisting of size-specific observations taken on successive census dates. For purposes of illustration, we applied our SSIPM method to a dataset collected in the kelp forests offshore of Pt. Lobos, California, USA, just south of Monterey Bay (36°31′1.56″ N, 121°56′33.36″ W). The Pt. Lobos region had a small marine protected area (MPA) in place since 1963, and this became a no-take MPA in 1973.

Table 1. Symbols used in this study.

Symbol	Definition
Integral projection	
$N(y,t)$	probability density of individuals of size y at time t
$K(y,x)$	IPM kernel giving probability density of transition from size x to y
$P(y,x)$	growth portion of the kernel k
$Q(y,x)$	reproductive portion of the kernel k
Ω	set of biologically reasonable sizes, x
h	IPM integration mesh resolution (bin width) (cm)
Demographic parameters	
L_∞	asymptotic maximum size (cm)
L_{CV}	coefficient of variation in size-at-age
k	von Bertalanffy growth rate (yr^{-1})
M	natural mortality rate (yr^{-1})
F	fishing mortality rate (yr^{-1})
F_1	F during the intial pre-data period (1990–1998; yr^{-1})
F_2	F for the period during which survey data were collected prior to MPA establishment (1989–2007; yr^{-1})
$\phi(x)$	fishing selectivity (probability that fish of size x are in the fishery)
μ_f	Mean size of entry to fishery (mean of $\phi(x)$; cm)
σ_f	standard deviation of size of entry to fishery (SD $\phi(x)$; cm)
$R(t)$	number of recruits in year t
$\rho(x)$	size distribution of age-0 recruits
Particle filter	
$N(t)$	true state of the process model
$X(t)$	observed data
ν	process error
σ_ν	standard deviation of process error term
ε	measurement error (variance)
σ_ε	standard deviation of measurement error term
$G(N,\nu)$	process model function
$H(N,\varepsilon)$	observation model function
$w_{i,t}$	weight (approximate likelihood) of particle i, time t
q	number of particles

Then, in 2007, that MPA was expanded in size as part of California's state-wide MPA implementation through the Marine Life Protection Act (MLPA; Kirlin et al. 2013, Botsford et al. 2014). The MLPA also mandates the monitoring and adaptive management of the newly implemented MPAs; i.e., assessment of whether fish populations increased as expected after protection inside MPAs. However, prior modeling studies have shown that in order to predict the expected increase of fished populations inside MPAs, one must know the level of fishing at that site prior to MPA establishment (White et al. 2011, 2013, Moffitt et al. 2013). This implies that setting expectations for management requires estimating the pre-2007 fishing rate at Pt. Lobos and other MPA sites. Estimating this unknown parameter was the motivation for our development of a SSIPM.

From 1999 to 2007 annual surveys of kelp forest fishes were made at seven study sites near Pt. Lobos using a sampling design that was spatially stratified within sites to capture the major gradients of environmental variability within the kelp forest (see Appendix S1 for details). For the model, we were interested in patterns of relative density among sites and years at the site scale rather than within-site gradients. Therefore we obtained an overall population density by summing the number of fish in each size class across all transects sampled within the site in a year. This sums across all of the major environmental gradients, giving a density that is comparable across sites and years. There was some variability in the number of transects used in some sites and years (Appendix S1), so we also tracked the number of transects used in each survey in order to correct for this during analysis.

We analyzed data from two of the most abundant species in the Pt. Lobos dataset: the blue rockfish (*Sebastes mystinus*) and the gopher rockfish (*S. carnatus*). Both are kelp forest residents and are commonly caught by recreational anglers; there is also some commercial harvest (Leet et al. 2002, Starr et al. 2002, 2010, Carr and Reed 2015). Typical home ranges of both species are less than 2 km^2 (Freiwald 2012, Green et al. 2014, Starr et al. 2015). Blue rockfish are found throughout the water column and prey primarily on salps and other zooplankton, while gopher rockfish are benthic and prey on crustaceans and small fish (Miller and Geibel 1973, Hallacher and Roberts 1985).

Uncertainty in the pre-MPA fishing mortality rate for these species arises, in part, from uncertainty regarding the local-scale realization of system-wide harvest regulations. Although both species have stock assessments that estimate harvest rates for the relevant time period, these estimates consisted of a single value for the entire central and northern California stock ranging from Pt. Conception to the Oregon border, a distance of >1,200 km (Key et al. 2005, 2008). However, human population density varies considerably along the coastline, and fishing effort can differ greatly over distances of a few km (Wilcox and Pomeroy 2003, Scholz et al. 2004). We therefore used the model to estimate the fishing rate as a free parameter, although we incorporated the existing information from the stock assessments by taking a Bayesian approach, using the stock-wide estimates of the harvest rate as a prior on the site-specific fishing rate we estimated (see *Decision point: prior distributions* for details).

In addition to pre-MPA fishing mortality, a second major source of uncertainty and variability is interannual variation in the recruitment of larval rockfish to the adult population. Both species have an annual peak of spawning (December–January for blue rockfish, March for gopher rockfish); the offspring spend 3–5 months in the plankton as larvae and post-larval juveniles before settling to nearshore kelp habitat in the late spring and early summer (Love et al. 2002). The offshore environment is highly variable in this region, and the survival of planktonic rockfish larvae is quite sensitive to variation in phytoplankton productivity, upwelling, sea surface temperature, and other factors at a range of

spatiotemporal scales. As a consequence, the size of the annual recruitment pulse entering the kelp forest varies by more than an order of magnitude, leading to boom and bust years that differ among rockfish species (Johnson 2006*a*, Caselle et al. 2010*a, b*). While annual recruitment is strongly correlated with oceanographic indices in Southern California, the Central California region is less predictable (Caselle et al. 2010*b*). Therefore, as we describe in *Decision point: open or closed demography*, we also estimated parameters that describe the magnitude of annual recruitment in each year.

Finally, we expect there to be process error in the form of annual variation in mortality and growth rates, as well as movements of fish in and out of the surveyed site, and measurement error in the visual census data themselves. Using a state-space model allows us to directly estimate and account for both of those sources of variability.

Constructing the IPM kernel

The first step in building an IPM is to construct the kernel using size-specific growth, mortality, and reproductive rates. This is done by constructing a size-growth function, $P(y,x)$, which gives the probability density function for growing (or shrinking) from each size x to y, multiplied by the probability of surviving the transition. Likewise, a fecundity function $Q(y,x)$ gives the density function for the number of offspring of size y produced by a parent of size x. The kernel is the sum of these two functions: $K(y,x) = P(y,x) + Q(y,x)$.

Decision point: parameter estimation.—The typical approach to IPM construction is to estimate the size-dependent parameters of P and Q from individual-scale data (Coulson 2012, Merow et al. 2014). Alternatively, Ghosh et al. (2012) and González and Martorell (2013) parameterized the kernel by finding the parameter values that afforded the best fit of the IPM to survey data. The latter approach can require estimating a large number of free parameters, which can lead to issues with parameter identifiability and the potential for biologically unrealistic parameter combinations (e.g., extremely low adult fecundity combined with extremely high juvenile survival; González and Martorell 2013). Therefore we took a hybrid approach: We took advantage of published literature estimates for some of the growth and mortality parameters, and then we used model fitting to estimate the remainder.

There is a wealth of recent literature describing the procedures for obtaining the size-growth and fecundity functions from empirical data (Coulson 2012, Metcalf et al. 2012, Merow et al. 2014, Rees et al. 2014), so we do not retread that ground here. For our implementation, size-specific demographic rates were already readily available from the stock assessments of the two species (Appendix S2; Key et al. 2005, 2008). For many fishes such information is available in the published literature (*available online;*[8] Froese and Pauly 2012), even if a stock assessment is not available. These life history variables are usually derived from studies with a broad geographic scale, and there can be spatial heterogeneity in some life history parameters at smaller scales (e.g., Caselle et al. 2011), so such estimates should be used with care.

Fish growth is typically indeterminate but asymptotic and is well described by the von Bertalanffy relationship, in which the mean growth Δx from t to $t + 1$ of a fish of length x at time t is $(L_\infty - x)e^{-k\Delta t}$, where L_∞ is the asymptotic maximum mean size, k is the growth rate, and Δt is the time step. Variability around the mean size is typically expressed as a constant coefficient of variation, because fish size at $t + 1$ varies proportionally to size at t. We thus modeled growth as a function in which the probability of growth to y from any particular size x followed a normal distribution, with the mean given by the von Bertalanffy relationship and standard deviation given by the coefficient of variation, L_{cv}, multiplied by that mean.

Regarding survival, we assumed that fish of all sizes experienced natural mortality at rate M (so the probability of survival over one year is e^{-M}), and we obtained estimates of that rate from the species' stock assessments. It is likely that mortality is actually size- and age-dependent, but in the absence of size- or age-specific survival data, we did not include that in our model. Additionally, fish outside MPAs are vulnerable to harvest in this system. Because fishing is often size-selective, we modeled size-specific survival as $e^{-(M+F\phi[x])}$, where F is the instantaneous fishing mortality rate (yr^{-1}) and $\phi(x)$ is the fishing selectivity function. The latter is a cumulative normal distribution with mean μ_f and standard deviation σ_f, so the effective harvest rate is zero for small sizes and begins to approach F as fish size approaches and exceeds μ_f, the minimum size retained in the fishery.

Decision point: mesh size, resolution, and integration.—Because it describes a continuous probability function over size, integration is required to obtain the number of individuals moving on to the next time step, $t + 1$. However, although it is mathematically defined as a continuous function, numerical calculations inevitably require that the kernel be discretized as a matrix. This requires decisions about the upper and lower limits of the mesh, the mesh resolution, and the numerical integration method. We describe the technical aspects of these points in Appendix S3.

Decision point: open or closed demography.—In most IPM implementations, the population being modeled is assumed to be closed, without any appreciable immigration or emigration and with reproduction modeled as a component of the kernel ($Q[y,x]$). However, in marine systems in particular, a given local population is likely to receive a large number of immigrants in the form of dispersing larvae that were spawned in other populations. For example, in the kelp forest fish species we modeled,

[8] www.FishBase.org

the planktonic larval stage lasts several months, during which time larvae can disperse far from their natal habitat. Because the spatial scale of our study sites is small relative to the presumed larval dispersal distance of rockfish, we assumed it was likely that most of the post-larval juveniles settling to the kelp forests in our dataset had been produced elsewhere. Therefore we structured our model as a completely open population, where all juveniles are considered to arrive as immigrants, and reproduction was not explicitly coupled to adult spawning within the local population. Thus we modified Eqs. 1 and 2 to include a term $R(t)$, the number of juvenile recruits arriving at time t, and we set the reproductive portion of the kernel $Q(y,x)$ to zero

$$N(y,t+1)=\int_{\Omega} K(y,x)N(x,t)dx+R(t)\rho(y)+\nu(y,t) \quad (3)$$

where $\rho(y)$ is the probability density function for initial recruit size. Of course, other model formulations are possible. Several studies with IPMs of marine species with dispersive larvae have compared models with closed (our Eqs. 1 and 2) and open dynamics (our Eq. 3) or a mix of both (Bruno et al. 2011, Madin et al. 2012, Yau et al. 2014). In general, modeling this type of species with closed dynamics without immigration would require a spatial scale large enough for the population to be largely self-seeding. In practice, it is possible to consider a mixed approach that includes both local reproduction and immigration (Yau et al. 2014), but a challenge would be to correctly estimate the probability that locally produced larvae return to the natal population (Burgess et al. 2014) and their mortality rate in the plankton (White et al. 2014), as well as any density-dependent post-settlement processes (e.g., Johnson 2006*a, b*). This is beyond the scope of our current modeling effort. In our model, recruitment, $R(t)$, is variable from year to year, reflecting the dynamic ocean environment of this system (as described previously in *Case study: a California marine protected area*).

Fitting the state-space IPM

We estimated the unknown parameters in the model (annual recruitment, harvest rate, and error term) by fitting the SSIPM to data. In principle, any kernel parameter could be estimated in this way, although of course in practice the uncertainty in the model fits will grow as we attempt to estimate more parameters.

We began model runs at the stable size distribution of a deterministic version of the IPM (see *Decision point: state-space filter*). However, the actual population was likely far from this size distribution at the beginning of observations in 1999. Therefore, we began the model ($t = 1$) at a distant year in the past (1990) and then ran the model forward, allowing variation in recruitment and process error during the nine years of burn-in prior to the first year of observations. This length of time is sufficient to ensure that the population would primarily consist of recruits from the post-1990 period. To reduce the overall number of model parameters, we assumed that the number of juvenile recruits $R(t)$ was constant for the 1990–1998 burn-in period but estimated a separate $R(t)$ for each of the nine years between 1999 and 2007 (when data were collected). We also assumed that fishing pressure may have differed in the 1990s and 2000s, so we estimated separate values of F for 1990–1998 (referred to as F_1) and 1999–2007 (referred to as F_2; this division is also consistent with an overall reduction in coast-wide fishing rates between the 1990s and the 2000s; Key et al. 2005, 2008). Additionally, we estimated the error term in the model. We assumed that the process error term $\nu(y,t)$ is normally distributed with mean 0 and standard deviation σ_ν, and we estimated the value of σ_ν (to avoid negative population density we constrained $N(y,t) > 0$). Because we assumed the survey data followed a Poisson distribution (see *Decision point: likelihood calculation*), there was not an explicit measurement error term (σ_ε) but had there been one could also have estimated it.

Two of the Pt. Lobos sites, Bluefish and Weston, were already inside a no-take MPA for the 1990–2007 period (this MPA was expanded in 2007 to include one other site, Monastery). Therefore we assumed $F_1 = F_2 = 0$ for those two sites, but still included data from those sites in the fitting process to estimate the recruitment and error parameters.

Regarding the recruitment rate $R(t)$, it is well known that kelp forest fishes, including our example species, experience high and potentially density-dependent mortality immediately following settlement from the plankton (Carr and Syms 2006, Johnson 2006*a, b*, 2007, White and Caselle 2008). However, that mortality is strongest and most density dependent within 1–2 months following settlement (late spring/early summer), whereas population surveys occurred in the late summer and autumn. Therefore the surveys only observe recruitment after much of the density-dependent mortality has acted. In our model we subsumed post-settlement, pre-census survival within $R(t)$, essentially modeling all pre-census factors as a single process, and assumed that a given recruit cohort's abundance was set at the time of census, and all subsequent mortality was density independent. Although there are options for modeling nonlinear demographic functions in an IPM context (Briggs et al. 2010), our approach is both simpler and reflects the limitations of the dataset.

We performed SSIPM model fits to the observed data for each of the two species. We assumed that because all the sites at Pt. Lobos were in close proximity, they received the same input of larval recruits each year and experienced the same fishing pressure. Thus we estimated single values of F_1, F_2, and $R(t)$ across all seven sites. We fit a separate process model to each site (allowing interannual process error to propagate within each site) but otherwise the collection of sites shared the same

parameters (except for the two sites where $F = 0$). The likelihood for the entire model was calculated as the product of the likelihoods for all sites in a given year. Not all sites were surveyed in all years, so likelihood calculations were only made for year-site combinations that had data.

Decision point: state-space filter.—A common approach to estimating $\boldsymbol{N}(y,t)$ given the observations $\boldsymbol{X}(y,t)$ is the Kalman filter (Kalman 1960, Meinhold and Singpurwalla 1983, Dennis et al. 2006), which provides an exact solution to the state-space problem for the special case of a linear model and normally distributed error terms. We used the alternative, less restrictive approach of a particle filter (Gordon et al. 1993, Knape and de Valpine 2012). Knape and de Valpine (2012) provide a thorough description of the particle filter procedure, but the basic approach is as follows (for convenience, we use matrix notation to refer to the continuous vector of actual abundance across all sizes at t, $\boldsymbol{N}(t)$ and the size distribution data vector, $\boldsymbol{X}(t)$):

1. For $t = 1$, simulate q independent random vectors $\boldsymbol{N}(t)_i'$ ($i = 1, 2, 3, \ldots. q$) from a distribution around $\boldsymbol{N}(t)$. This set of vectors $\boldsymbol{N}(t)_i'$ are the particles, and they represent the initial distribution of possible true hidden states.
2. Calculate weights $\boldsymbol{w}_{t,i}$ for each particle that indicate how well each particle matches the corresponding observed data, $\boldsymbol{X}(t)$. The weight is essentially a likelihood: $w_{i,t} = L(\boldsymbol{X}[t]|H[\boldsymbol{N}(t)_i', \varepsilon(t)])$. (Recall that H was defined previously as the function translating the actual size distribution into observed sizes.)
3. Resample the particles with replacement according to their relative weights,

$$\frac{w_{t,i}}{\sum_{i=1}^{q} w_{t,i}}, \tag{4}$$

   to obtain an updated estimate of the distribution of hidden states $\boldsymbol{N}(t)_i$ based on comparison to the data.

4. Advance the model to obtain a new set of particles: $\boldsymbol{N}(t+1)_i' = G(\boldsymbol{N}[t]_i, \nu[t])$.

Steps 1–4 are then repeated for each step of the time series. At each step, the resampled particles, $\boldsymbol{N}(t)_i$, represent the approximate distribution of the hidden state at t given the observations up to that point, $\boldsymbol{X}(1{:}t)$, and the mean of those resampled particles can be taken as a representation of $\boldsymbol{N}(t)$. Helpfully, the weights also approximate the likelihood of the data given the current parameter values, θ,

$$L(X(1{:}T)|\theta) = \prod_{t=1}^{T} \frac{1}{q \sum_{i=1}^{q} w_{t,i}}. \tag{5}$$

In order to start the filter algorithm, it is necessary to specify the true state at $t = 1$. One option is simply to use the stationary distribution of the process model (i.e., the stable size distribution of the deterministic IPM) as a starting point (de Valpine and Hastings 2002). We took this approach and simulated the initial particles $N(1)_i$ by adding process noise ν to the initial distribution $N(1)$ (where $t = 1$ corresponds to the year 1990). We used the mean of the prior distribution on σ_ν to simulate ν for that purpose.

Decision point: likelihood calculation.—The calculation of the particle weights $w_{i,t}$ requires an expression for the likelihood $L(\boldsymbol{X}[t]|H[\boldsymbol{N}(t), \varepsilon(t)])$. The form of the likelihood largely depends on the form of the observational data. For example, integer count data of individual organisms (such as the counts of individual fish of particular length in the Pt. Lobos dataset) usually follow a Poisson distribution. Therefore we used a Poisson likelihood with expectation

$$\langle X(x,t) \rangle = \int_{x-h\frac{1}{2}}^{x+h\frac{1}{2}} H(N[x,t])dx \tag{6}$$

for each width-h interval of the integration mesh. No explicit measurement error term is required because the variance equals the mean in the Poisson distribution. Recall that the number of transects sampled varied among years and sites, and the data $\boldsymbol{X}(t)$ were sums across all transects. Therefore our function H includes multiplying $\boldsymbol{N}(t)$ by the number of transects sampled in that year in order to obtain the Poisson expectation for the number of fish observed. Other likelihood functions may be more appropriate for datasets with different error structures or sampled using different methods.

Decision point: observation ogive.—It is possible that not all sizes of the organism are observed similarly within the surveyed habitat, and this must be accounted for in the function $H(\boldsymbol{N}(t))$. For example, in blue rockfish, there is evidence that larger individuals migrate offshore, out of the visual census area within the kelp forest. Starr et al. (2015) conducted a fishery-independent hook-and-line survey just offshore of the Pt. Lobos sites where underwater visual censuses were conducted during the same time period. The data from those hook and line surveys revealed that the size distribution of blue rockfish shifted to larger individuals offshore of the kelp forests, suggesting that individuals moved offshore as they grew larger (Appendix S4). The IPM does not explicitly account for this movement, so it would predict a size distribution of the blue rockfish population that is more weighted towards larger size classes than would be described by the visual surveys. That IPM-predicted distribution would be more relevant to what is susceptible to the fishery, but we must account for the difference when comparing the model to visual survey data. Therefore we created an ogive function giving the probability of a fish being observed in the nearshore surveys. We did this by comparing the combined size distributions across all sites and years in the kelp

forest surveys and the hook-and-line surveys and estimated a cumulative normal distribution function that described the probability of being present in the hook-and-line survey. This function had a mean equal to the mean size of the hook-and-line data and a standard deviation that gave a <1% probability of fish <10 cm being observed in the hook-and-line data (Appendix S4). One minus that probability is then the probability of a size class being present in the visual survey dataset, and we used that probability to specify the function translating the hidden state into the observed data, $H(\boldsymbol{N}[t])$. For the other species we studied, gopher rockfish, there is no evidence for ontogenetic offshore migration, so the function H is simply the identity function (albeit multiplied by the number of transects, as described earlier). In other cases it would be possible to represent the difficulty of observing very small juvenile stages or similar observational difficulties using the same procedure.

Decision point: number of particles.—In general, the accuracy and precision of the particle filter method (i.e., its ability to consistently reproduce correct results) improves if more particles are used. This is because a larger number of particles will provide a better representation of the hidden process state, $\boldsymbol{N}(t)$. However, there are diminishing returns and the computational demand increases linearly with the number of particles. We determined the optimal number of particles to simulate by calculating the coefficient of variation (CV) of the likelihood (Eq. 4) for 100 independent simulations of the IPM for a range of the number of particles, q (Appendix S5). The logic of this is that once there are a sufficient number of particles, the IPM produces a consistent estimate of the likelihood for the same set of parameter values, which is thus suitable for likelihood-based parameter estimation. For our IPM, we found that the CV of likelihood had an elbow near $q = 100$, at which point the CV began to decrease much more slowly with increasing q (Appendix S5). Therefore we used $q = 100$ particles in our model.

Parameter estimation: Markov chain Monte Carlo

Given the high dimensionality of the unknown parameter space (e.g., a distinct value of $R(t)$ for each model year) and the possibility of a multi-modal likelihood surface, the most practical choice for parameter estimation in our implementation was Markov chain Monte Carlo (MCMC; as in Knape and de Valpine 2012). In cases where it is possible to determine an expression for the likelihood of each parameter conditional on the others, a Gibbs sampler (Casella and George 1992) using software such as BUGS would be an efficient MCMC choice (*available online*).[9] We were unable to find expressions for the marginal likelihoods of each parameter, so instead we used a delayed-rejection one-at-a-time Metropolis–Hastings algorithm (Chib and Greenberg 1995, Green and Mira 2001).

[9] http://www.mrc-bsu.cam.ac.uk/software/bugs/

The details of implementing Metropolis–Hastings MCMC are described extensively elsewhere (e.g., Brooks et al. 2011). In our implementation, we followed Knape and de Valpine (2012) in using the particle-filter approximation of the likelihood $L(\boldsymbol{X}[1{:}T]|\theta)$ (Eq. 5) for the Metropolis–Hastings step. We updated the candidate parameters F_1, F_2, $R(t)$, and σ_v one at a time. For F_1, F_2, and $R(t)$, the proposal distributions were normal distributions centered at the current state of the Markov chain and with a coefficient of variation that decreased geometrically during the delayed-rejection process. For σ_v, we used an inverse gamma proposal distribution with a mean equal to the current parameter state and a shape parameter that decreased geometrically during delayed rejection. MCMC runs were made with 10^4 steps (this number is relatively small because the delayed-rejection procedure allowed rapid mixing of the chain), with the first 5×10^3 steps discarded as burn-in. We ran three chains with random initial parameter distributions for each MCMC instance, checked chains for convergence using the scale reduction factor diagnostic (Gelman and Shirley 2011), and pooled chains to estimate the posterior distribution of each parameter.

Decision point: prior distributions.—Like all Bayesian methods, the MCMC procedure requires a prior distribution for each parameter. If no prior information is actually available, then it is reasonable to choose an uninformative prior. In our case, we had prior estimates of F for the 1990–1998 and 1999–2007 time period because both species had recent stock assessments that estimated the harvest rate (Key et al. 2005, 2008). These were stock-wide estimates of the harvest rate; that is, harvest was estimated for the entire California population of each species, with the population assumed to be well-mixed at that spatial scale. We expect that the actual harvest rate at Pt. Lobos was likely to be rather different from that coast-wide estimate, so we developed somewhat weak priors based on those estimates (Table 2). Both stock assessments reported estimates of F for each year, so we created prior distributions with mean equal to be the mean value of F in each data range (1990–1998, 1999–2007) and standard deviation equal to the standard deviation of F in each range. For the recruitment rate and the process error term, we did not have any prior information, so we used uninformative priors (Table 2).

Simulated data

In order to test the accuracy of parameter estimation by the SSIPM, we also simulated blue rockfish datasets with specified values of F_1, F_2, and $R(t)$ and fit the model to those simulated data. The simulated datasets were the same length as the real data (9 yr). To avoid circularity, the simulated data were not created using the same IPM we used for fitting; instead we simulated the dynamics of an age-structured model (the same model used by Moffitt et al. 2013) using the same demographic parameters as

TABLE 2. Prior distributions used in MCMC estimation.

Parameter	Blue rockfish	Gopher rockfish	Simulated data
F_1	lognormal ($\mu = e^{0.25}$, $\sigma = 0.23$)†	lognormal ($\mu = e^{0.24}$, $\sigma = 0.44$)#	lognormal ($\mu = 1 \times 10^{-5}$, $\sigma = 10$)
F_2	lognormal ($\mu = e^{0.094}$, $\sigma = 0.376$)	lognormal ($\mu = e^{0.10}$, $\sigma = 0.44$)#	lognormal ($\mu = 1 \times 10^{-5}$, $\sigma = 10$)
log $R(t)$	Normal ($\mu = 0.87$, $\sigma = 1.70$)‡	Normal ($\mu = -2.22$, $\sigma = 1.25$)‡	Normal ($\mu = 0.87$, $\sigma = 1.70$)‡
σ_v	Inv. gamma (2,10)§	Inv. gamma (2,10)§	Inv. gamma (2,10)§

† Based on baseline estimates in (Key et al. 2008).
‡ Estimated mean annual recruitment in 1999–2007 in the survey data.
§ Relatively flat prior with expectation of 0.1.
# Based on baseline estimates in (Key et al. 2005).

the blue rockfish IPM (Appendix S2). We added process error to the model by making the annual mortality rate a random normal variable with mean M (from Table 1) and standard deviation 0.01. We then converted age distributions into size distributions using the von Bertalanffy relationship and then sampled the data by binning abundances into bins of width 3 cm, which approximates the accuracy of the diver-collected field data (in principle, fish are sized to the nearest cm but in practice there are more observations of fish at round number lengths such as 20 cm instead of 19 or 21 cm). We fit the state-space IPM to three sets of simulated data with $F_1 = F_2 = 0$, 0.05, or 0.1/yr, respectively, to gauge how well the model could estimate F_2 at different intensities. There were 10 replicate simulated datasets in each of these three sets. We also created three additional sets of simulated data (again, 10 replicate datasets each) to investigate the robustness of the model to data of lower quality than those we used. These three sets had (1) only seven years of data rather than nine, (2) only three years of data, and (3) nine years of data, but fish lengths were binned into 10-cm bins, approximating a much coarser scale, less accurate survey technique (such as might be available from citizen science surveys). For these model fits we used uninformative priors for all parameters (Table 2).

Software

All simulations and model fitting were performed in Matlab 8.5.0 (R2015a; Mathworks, Natick, Massachusetts, USA). All model code and example data are available online.[10]

RESULTS

Simulated data

The model fits to simulated blue rockfish survey data were able to estimate the actual values of the unknown parameters F_1, F_2, and $R(t)$ with reasonable accuracy (Figs. 2 and 3). The model had greatest accuracy for the moderate value of F_2 ($F = 0.05$/yr); the model estimates deviated from the actual value of F_2 the least (the mean estimate across all 10 simulated datasets was 0.0527/yr), and all simulations produced posterior distributions with a 95% credible interval (CI) that contained the actual value of F_2 (Fig. 2b). When the actual value of F_2 was higher (0.1/yr), the model still produced posterior CIs that included the true value in all but one case, and the CIs were sometimes wider indicating reduced precision (Fig. 2c; the mean estimate was 0.1027/yr). Simulations with no fishing ($F_2 = 0$) posed a unique problem for the model because negative values of F are not possible. Consequently, the posterior distribution of F_2 was necessarily asymmetrical and the mean was biased towards positive values (although the mode was always at the extreme left edge of the distribution; *data not shown*). Nonetheless, the model did estimate very low values of F_2 (mean: 0.0024/yr), and the 95% CI were narrower compared to those of the other actual F_2 values (Fig. 2a).

Reductions in data quality had a serious effect on the model's accuracy in estimating F_2. When there were only seven years of simulated data, the posterior mean estimates of F_2 bracketed the actual value of 0.05/yr, and the 95% credible intervals included the actual value for only four of the 10 simulations (Fig. 2e). When there were only three years of simulated data (Fig. 2f), or length data had coarser resolution, with sizes binned into 10 cm intervals (Fig. 2g), the model estimates of F_2 were biased towards much lower values (or in some cases unreasonably high values; Fig. 2g), and the 95% CIs almost never included the actual value.

Regardless of the value of F, the models provided accurate estimates of the number of annual recruits, $R(t)$, with estimated values close to the actual value (Fig. 3a). The model also produced consistent estimates of process error (σ_v, Fig. 3b; two simulations had fitted values of σ_v that were much higher than the rest, on the order of 0.1). Estimates of σ_v are not directly comparable to actual values because in the simulated datasets, process error was introduced as variation in natural mortality, M, rather than in the absolute numbers of fish as the particle filter estimates.

The model fit the simulated blue rockfish size distributions well (Fig. 4). The effects of increasing F were clearly apparent in the IPM fits as an increasingly truncated right-hand tail of the size distribution (compare Fig. 4a–c). The ability of the model to detect these

[10] https://github.com/jwilsonwhite/IPM_statespace (http://dx.doi.org/10.5281/zenodo/56574)

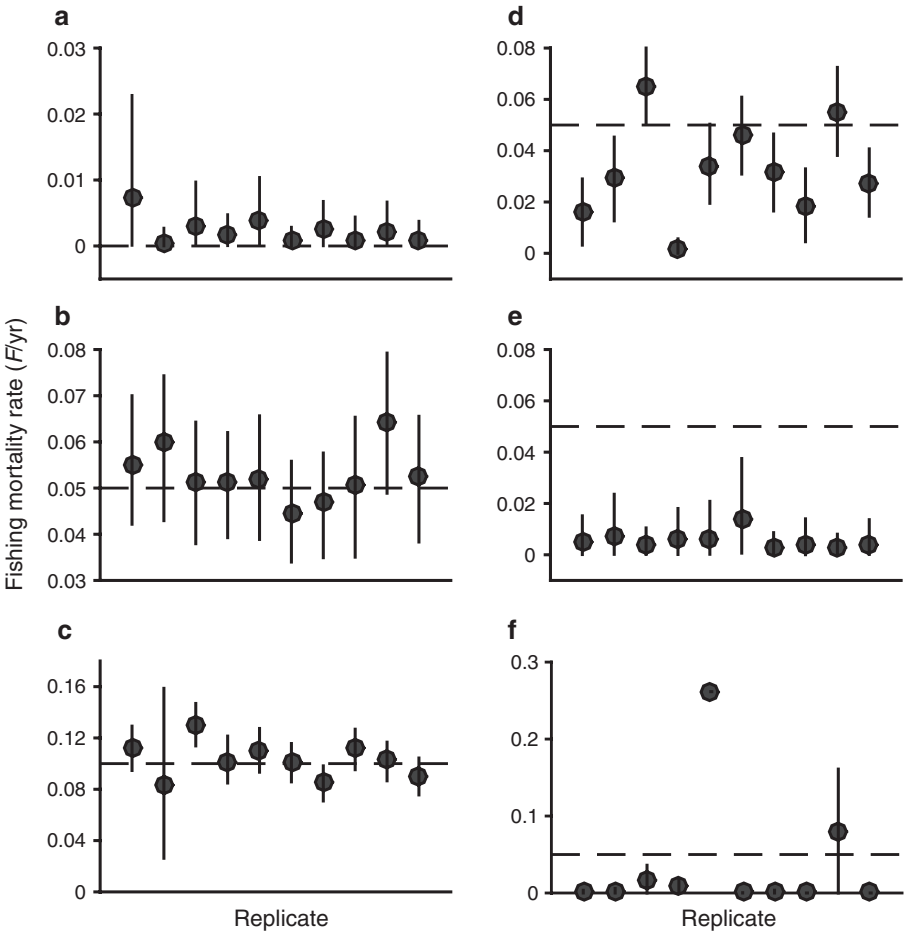

FIG. 2. Posterior estimates of the fishing mortality rate for model fits to simulated data. Simulated datasets were 18 yr long, but only the final 9 yr were used to fit the state-space IPM. The posterior estimate of F_2 (yr^{-1}, circles with bars) for the final 9 yr are shown for 10 replicate datasets with actual F_2 (dashed line) equal to (a) 0/yr, (b) 0.05/yr, and (c) 0.1/yr. Additionally, the model was fit to replicate datasets with actual $F_2 = 0.05$/yr and (d) only 7 yr of data, (e) only 3 yr of data, and (f) 9 yr of data but observations binned into 10 cm length bins. Markers indicate posterior mean, lines indicate 95% posterior credible intervals. The credible interval lies entirely within the diameter of the marker for some points.

differences in the largest size classes is key to accurate estimation of F. The model did slightly underestimate overall density, despite capturing differences in overall shape. This is likely a consequence of the binning artifact introduced into the simulated dataset: fish of length 19, 20, and 21 cm were all pooled into the 20 cm bin (for example), so the model contended with data that had a peaky distribution with zeros in many bins. Those zeros would tend to depress the overall density estimate somewhat. In Fig. 4, we chose to represent that data distribution as a density histogram with 3-cm-wide bins in order to facilitate comparison with the continuous size distribution in the model.

Case study: a California marine protected area

The model also produced reasonable fits to the size distribution data for both blue rockfish (Fig. 5) and gopher rockfish (Fig. 6) at Pt. Lobos. Recall that each individual study site within the Pt. Lobos area was assumed to have the same annual recruitment and fishing rate (except for the two sites inside MPAs, with $F = 0$), though each site had independent process errors. Consequently, the predicted population distribution is quite similar across sites, leading to an overall good fit even though there are some deviations in particular years for some sites (e.g., the model overestimated abundance of blue rockfish at site Bluefish in 2007; Fig. 5e). In particular, the model accurately captured the distinct differences in size distribution between years that had a strong pulse of juvenile blue rockfish recruits (1999; Fig. 5a, b) and those that did not (2006; Fig. 5c, d). The match between model fit and observed data was much better for blue rockfish than for gopher rockfish because the former species was more abundant, so the data had a more continuous size distribution. Similar to blue rockfish, the model identified a strong recruitment year for gopher rockfish observed by diver surveys in

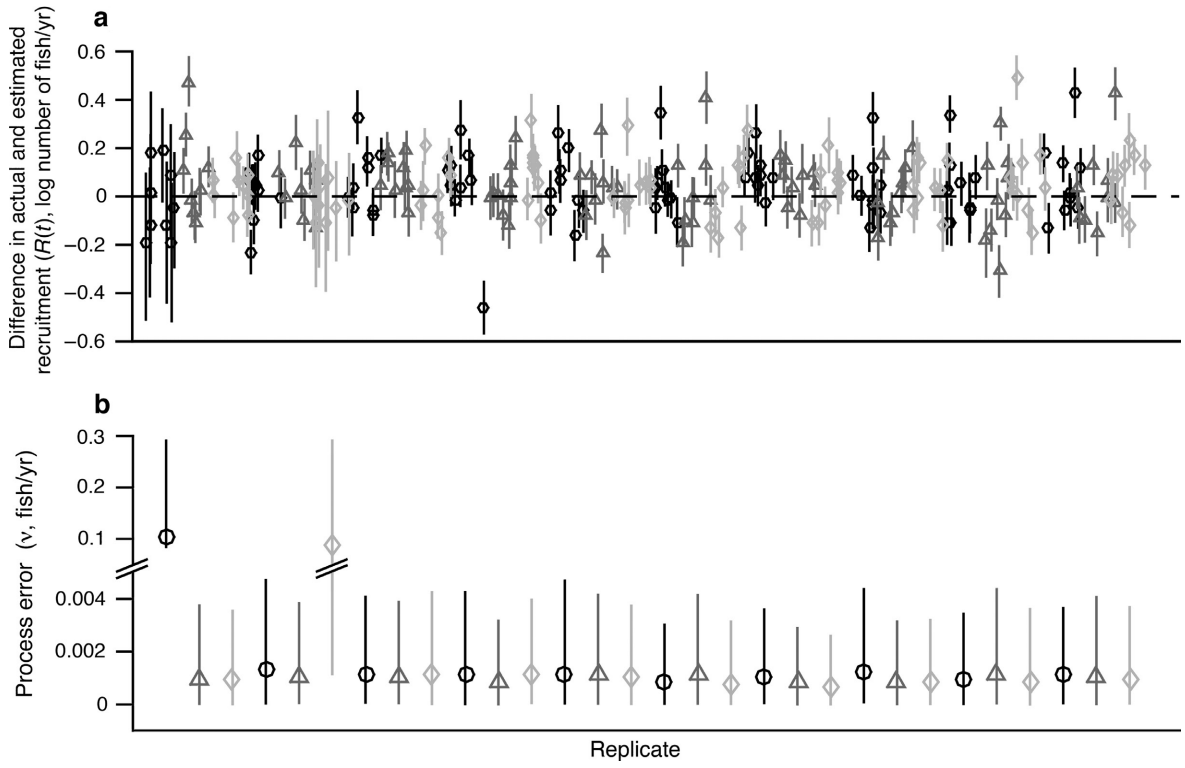

FIG. 3. Posterior estimates of larval recruitment and process error for model fits to simulated data. Panels show posterior estimates of (a) the annual larval recruitment rate, $R(t)$ (units, log number of fish/yr) and (b) process error ν, from state-space IPM fits to simulated data with different values of the fishing mortality rate, F_2, and annual variation in recruitment. In (a), posteriors are shown for each of the 9 yr of data for each of 10 replicate datasets; values are expressed as the difference between the actual value and the posterior estimate. In (b), raw posterior values are shown. Results are shown for actual values of F_2 equal to 0/yr (circles), 0.05/yr (triangles), and 0.1/yr (diamonds; the same datasets for which fits are shown in Fig. 2a–c). Markers indicate posterior mean, lines indicate 95% posterior credible interval. Note the break in the vertical axis scale in (b).

1999 (Fig. 6a, b), but unlike blue rockfish the model again predicted strong recruitment in subsequent years for gopher rockfish (2006, 2007; Fig. 6c–f) when very few recruits were actually observed.

The model produced similar posterior estimates of F_2 in the two species. The posterior mean was 0.19/yr for both blue rockfish (Fig. 7a) and gopher rockfish (Fig. 7b). Both estimates were >2 times higher than the prior estimate derived from the stock assessment (Fig. 7). The posterior distribution for blue rockfish was much narrower, indicating higher confidence in the value of F_2. That species was more common in the survey data, allowing a better model fit and higher confidence in parameter estimates (typically >200 blue rockfish were observed per site in a given year vs. <30 gopher rockfish).

As in the simulated data, the effects of fishing on the predicted distribution can be observed by comparing the fitted distribution for the site Bluefish to that for the site Monastery; the former was in a no-take MPA during the 1999–2007 period and so was assumed to have no fishing during this time, and the latter site became an MPA in 2007. This is reflected in the slightly thicker right-hand tail of the distribution, indicating a higher probability of observing large fish (Figs. 5 and 6).

DISCUSSION

We have described a method of implementing an integral projection model (IPM) in a state-space context. This allows ecologists to take advantage of the powerful attributes of IPMs, such as the small parameter space compared to structured matrix models, when fitting process models to data to estimate unknown parameters. This is particularly useful for the wealth of size-structured survey data available for a variety of systems. While IPMs are rapidly gaining popularity for prospective analysis of population growth rates and other demographic statistics (Coulson 2012), our state-space approach allows us to apply an IPM retrospectively, using time series data to estimate unknown parameter values as well as the underlying process-model state (as opposed to merely the observed data) at the present time. The SSIPM approach advances recent efforts to fit IPMs to size-structured data (González and Martorell 2013) by explicitly accounting for process and measurement error. It is also distinct from similar efforts to estimate a key unknown value, fishing mortality rates in a stock-assessment context (e.g., Key et al. 2008) because it is a purely size-based approach and does not require

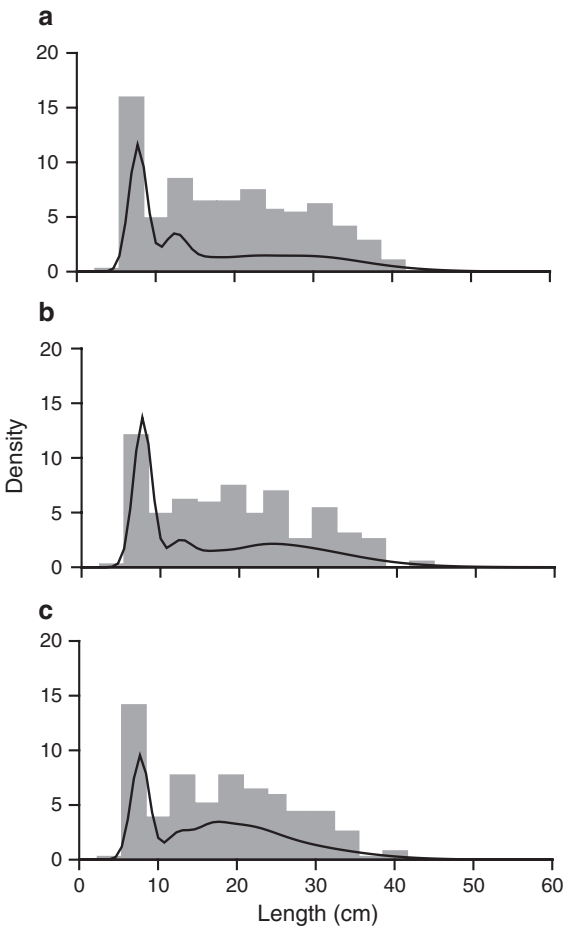

FIG. 4. Fits of the state-space IPM to simulated data. The IPM (curves) was fit to simulated length–abundance data (histogram) with actual values of the fishing mortality rate, F, equal to (a) 0/yr, (b) 0.05/yr, and (c) 0.1/yr. Fits are shown for one representative replicate dataset, for the final year of simulated data. Both model and data are presented as density distributions; actual abundances are obtained by integrating over a length interval (model) or multiplying by bin width (data).

difficult-to-obtain age data (aging fish requires killing them) as typical joint length/age-structured stock assessments do (Methot and Wetzel 2013).

Our analysis of simulated data revealed both the power of state-space estimation and some limitations of the method. In general, the model was able to estimate an accurate fishing mortality rate from noisy size-structured data; nearly all of the posterior estimates of F (for datasets with nine years of observations) had confidence regions that included the true simulated value (Fig. 2). This is not a trivial estimation problem because only the right-hand tail of the size distribution is truncated by fishing. The estimates had better accuracy for smaller values of F (for $F > 0$), which may be because as the fishing rate increases, fewer fish survive to large sizes and it is more difficult for the model to detect subtle changes in the tail of the size distribution. The case in which the true value of $F = 0$ presents a different problem; because F cannot, by definition, be negative, the true value lies at the boundary of the supported region of parameter space, so the posterior mean estimate will necessarily be greater than the true value (because the tail of the posterior distribution extends only into positive values). Nonetheless, the model estimated very small values of F (one could argue they are effectively equivalent to zero from a demographic standpoint) with quite narrow posterior 95% CIs. Therefore it would be possible to make a strong inference that the fishing mortality rate is not greater than the upper bound of the posterior 95% CI region.

The model also estimated the larval recruitment rate with high accuracy and precision. Estimation of that quantity in our model was less challenging: while estimating F depends on the shape of the right hand of the size distribution, estimating $R(t)$ merely depends on the integral of recruit size classes, which typically form a distinct peak in the size distribution. The model's ability to detect boom years with large recruitment pulses vs. bust years of recruitment failure was particularly evident in the blue rockfish dataset (Fig. 5). Accurate estimation of this process is an important advantage of the model, because larval recruitment is often a highly uncertain parameter in marine systems.

Our success in estimating harvest and recruitment rates belies some potential challenges in SSIPM estimation. First is the issue of parameter identifiability. If we had attempted to estimate all of the kernel parameters from the size distribution data, we may have obtained biologically unrealistic parameter combinations. For example, it is conceivable that a truncated size distribution could be fit equally by a model with high F (as in our examples) or with $F = 0$ and a reduced adult growth rate. González and Martorell (2013) encountered this issue and addressed it ad hoc by discarding a subset of maximum likelihood parameter estimates. Ghosh et al. (2012) also had some difficulty with parameter identifiability and so held some parameters and hyperparameters at arbitrary constant values in their hierarchical model. We avoided this problem in three ways. First, we limited the number of parameters to be estimated by using growth parameters from the literature. Second, the parameters to be estimated were length-specific and operated on distinct regions of the size distribution (new recruits and larger fished individuals); estimating both recruitment and juvenile mortality (for example) would be more challenging. Finally, we chose Bayesian priors that shifted the posterior away from implausible values. Using these strategies, it should be possible to estimate more than the two parameters we did here, such as environment-dependent hyperparameters for recruitment (e.g., allowing recruitment to depend on sea surface temperature or upwelling indices) or more complex size-dependent mortality rates.

A second, related challenge is that IPMs are vulnerable to an ecological fallacy, in which statistical patterns estimated at one scale are erroneously used to make

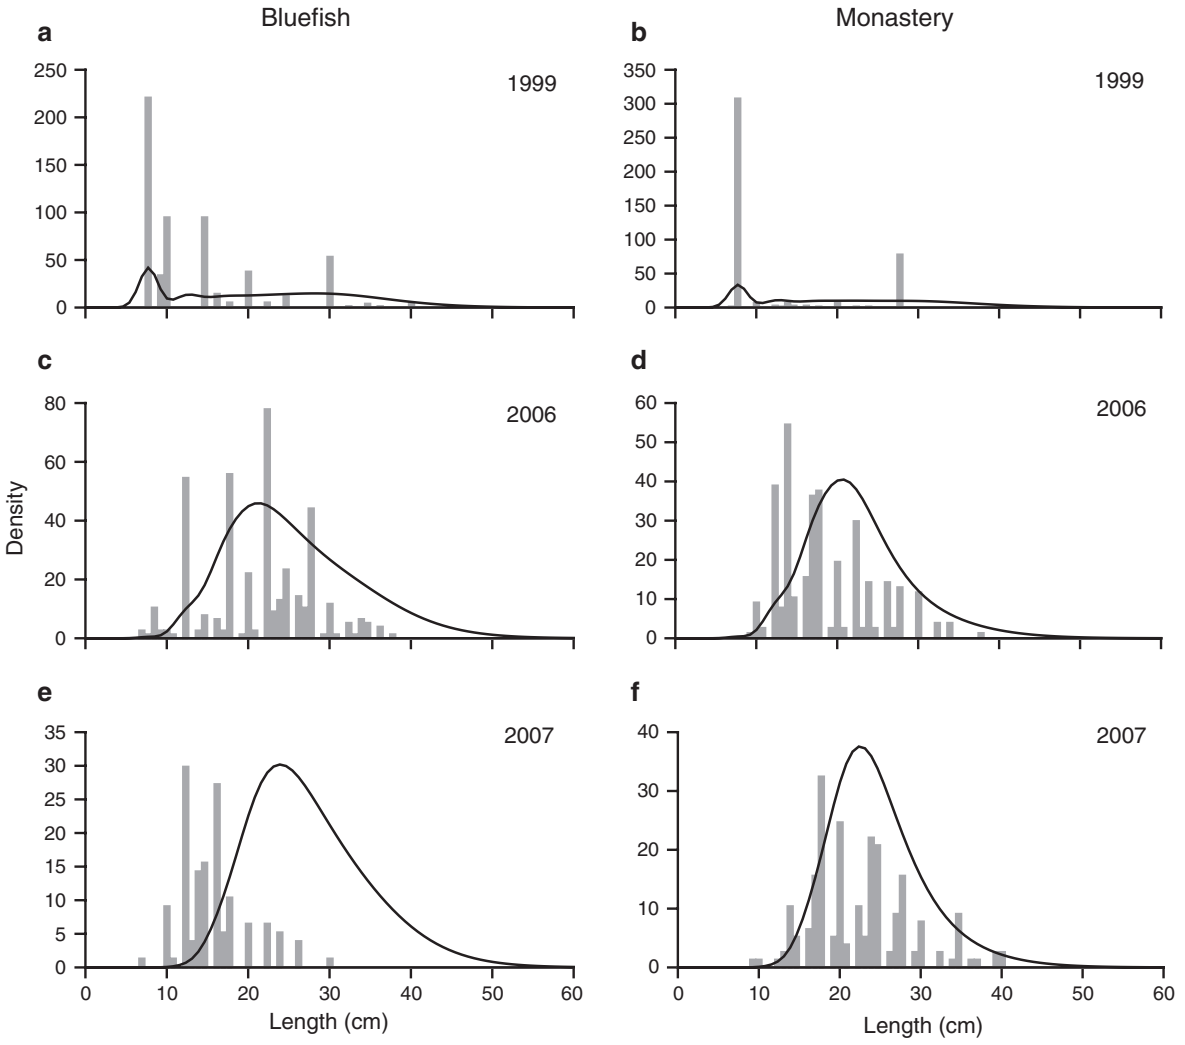

FIG. 5. IPM fit to blue rockfish (*Sebastes mystinus*) data. Gray histograms are the size distributions of blue rockfish observed at two sites in the Pt. Lobos region, (a, c, e) Bluefish and (b, d, f) Monastery in three representative years of the 1999–2007 dataset, (a, b) 1999, (c, d) 2006, and (e, f) 2007. State-space IPM fits to the data are displayed in black. Bluefish was inside a state marine reserve prior to 2007, so no fishing was assumed to occur there. Both model and data are presented as density distributions; actual abundances are obtained by integrating over a length interval (model) or multiplying by bar width (data).

inferences at a different scale (Clark 2003, Clark et al. 2011). Traditional IPM analyses are vulnerable to this error, because individual-scale relationships are used to project population-scale size distributions (Ghosh et al. 2012). The state-space approach avoids this by estimating kernel parameters directly at the population scale (Ghosh et al. 2012). Our SSIPM implementation is somewhat vulnerable to the fallacy because we used individual-scale age-size estimates to parameterize the growth kernel; this would be of particular concern if there were evidence of covariation in growth and mortality risk or perhaps spatial covariation in growth and harvest (e.g., fishery-induced selection on growth; Hutchings and Fraser 2008). Our growth estimates were derived from coast-wide data that spanned spatial gradients in harvest effort and should be less vulnerable to that error. Nonetheless, it is important to avoid the alternative ecological fallacy of inferring individual-scale relationships from population-scale SSIPM fits, such as assuming that fish in all regions of the kelp forest have equivalent predation risk.

The accuracy of the posterior estimates depend greatly on the quality of the data. The California MPA dataset (and the simulated datasets designed to mimic it) are of unusual quality in terms of the length of the time series (for a marine system), amount of replication, and precision of size estimates (to the nearest cm). In particular, such precise size–abundance estimates are very difficult to obtain from most underwater visual censuses, although this is perhaps less of a limitation for surveys of sessile benthic organisms or terrestrial plants.

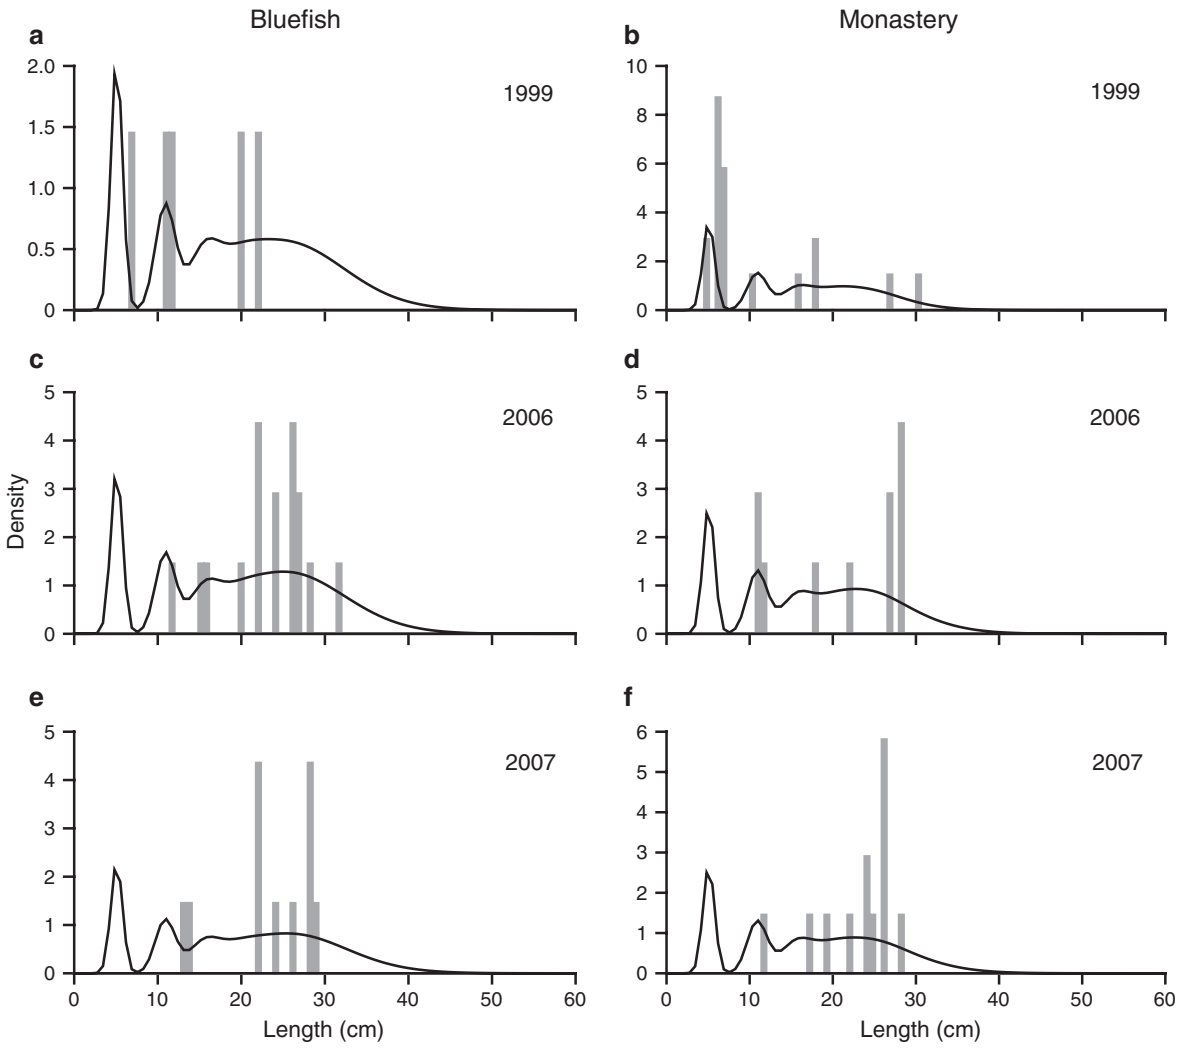

FIG. 6. IPM fit to gopher rockfish (*Sebastes carnatus*) data. Gray histograms are the size distributions of gopher rockfish observed at two sites in the Pt. Lobos region, (a, c, e) Bluefish and (b, d, f) Monastery in three representative years of the 1999–2007 dataset, (a, b) 1999, (c, d) 2006, and (e, f) 2007. State-space IPM fits to the data are displayed in black. Bluefish was inside a state marine reserve prior to 2007, so no fishing was assumed to occur there. Both model and data are presented as density distributions; actual abundances are obtained by integrating over a length interval (model) or multiplying by bar width (data).

When we simulated lower-quality data, we found that both the accuracy and precision of posterior parameter estimates were impaired (even for datasets with only seven years of data vs. the nine years in the full datasets). This was particularly true when the length data were binned much more coarsely in the dataset. Additionally, it is important to have enough observations in the sample to accurately characterize the shape of the size distribution. In our California example, there were an order of magnitude more individual blue rockfish than gopher rockfish in the survey data. Consequently the gopher rockfish size distributions were quite sparse, and the posterior estimates of the harvest rate were less precise, with considerably wider credible intervals (Fig. 7). Together, these caveats suggest that while our approach could be applied with some success to less robust datasets, additional simulated data analyses such as these should be used to characterize the expected accuracy of any implementation. Furthermore, it is clear that the accuracy of model projections to inform expectations for the rate of population increase in MPAs is strongly influenced by the quality of monitoring programs (i.e., frequency of surveys, resolution of size estimates, and choice of appropriate spatiotemporal scales; Carr et al. 2011, Cvitanovic et al. 2013). High-quality, long-term monitoring surveys are costly but allow inferences that are not possible using occasional snapshot surveys (e.g., Babcock et al. 2010).

In principle, this approach does not actually require full size-structured distributions. The process model

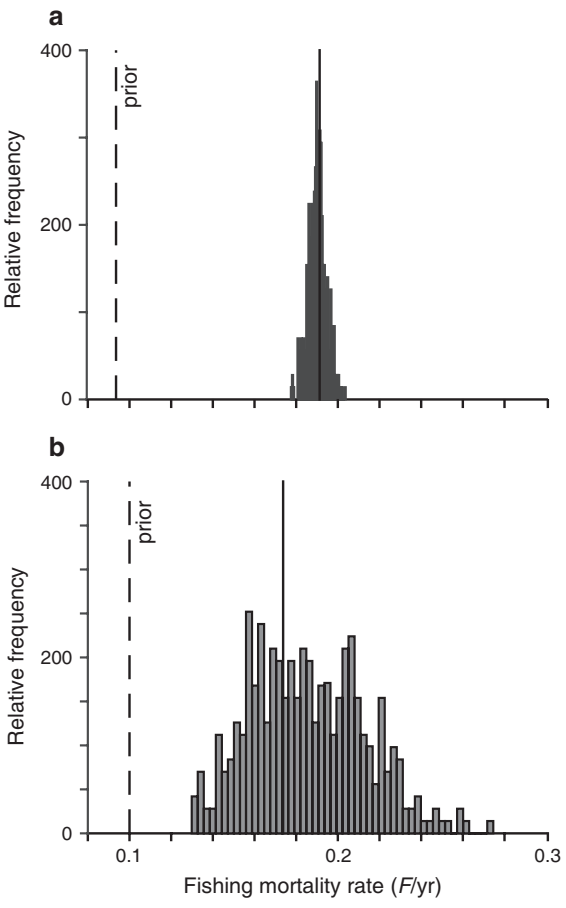

FIG. 7. Posterior distributions of the fishing mortality rate for blue rockfish and gopher rockfish. The posterior distribution of fishing mortality rate, F_2 (yr^{-1}), for 1999–2007 was estimated from the state-space IPM fit to data for (a) blue rockfish (*Sebastes mystinus*) and (b) gopher rockfish (*Sebastes carnatus*) at sites near Pt. Lobos, California. The dashed vertical line indicates the mean of the prior distribution on F based on coastwide stock assessments; the solid vertical line indicates the mean of the posterior distribution.

could be a size-structured IPM, and the observation model $H(N)$ could simply calculate the likelihood of observing a particular mean abundance, mean size, mean biomass, etc., given a particular size-structured state. We did not simulate this type of usage, but it would likely have much reduced accuracy and precision because of the loss of information when the size distribution is summarized with a single statistic. This would particularly hamper estimation of an unknown parameter such as F because of the loss of information on the right-hand tail of the size distribution.

Case study: California MPAs

In the California dataset, we estimated fishing mortality rates that were substantially different from the prior estimate from the stock assessment. It is rare to obtain such fine-scale, site-specific estimates of F in a fisheries study, and the fact that the estimated values differed from the coast-wide stock assessment estimate heightens their importance. In particular, we estimated a value of F nearly double the prior estimate for both species. That the posterior values were higher than the prior in both cases highlights the pitfall of applying a stock-assessment-based estimate of F to an individual MPA. The stock assessments (Key et al. 2005, 2008) estimate F over a large geographic extent, averaging over locations with varying levels of fishing effort. Point Lobos is easily accessible by boat by fishermen in Monterey Bay, and thus would be expected to have higher fishing effort than the coast-wide average.

Because the SSIPM estimates the current, hidden state of the underlying process model (as opposed to merely the observed state), it could be advantageous for making short-term projections for adaptive management. This will be particularly valuable in cases where the size distribution is far from its stable state and strongly affected by recent events, such as pulses or droughts in larval recruitment. By estimating the fishing rate and the actual size distribution at the time of MPA implementation, it is possible to make informed predictions about the likely trajectory of the population in the next decade, which may deviate greatly from the long-term asymptotic growth rate expected within an MPA (White et al. 2013). Those model predictions (or perhaps a suite of predictions reflecting different possible future environmental or management conditions, e.g., White et al. 2010) could later be compared to future monitoring data to determine if the population is following the expected trajectory. The prediction-comparison step is key to assessing management efficacy in adaptive management (Walters 1986) but is not possible with a purely asymptotic analysis of a traditional IPM. This analysis is beyond the scope of the current study, but elsewhere (K. J. Nickols et al., *unpublished manuscript*) we have applied our SSIPM approach to a broader suite of kelp forest fishes in Pt. Lobos and two other nearby MPAs, estimated the pre-MPA fishing rates, and projected the expected trajectories. We refer the reader to that paper for a fuller exploration of the details and consequences of these patterns for MPA management in California.

Future directions

We have outlined a relatively basic IPM structure here, with a case study tailored to the specific ecological context we were modeling. Future efforts could build on this framework to provide more realistic representations of ecological processes, better fits to data, and stronger tests of ecological hypotheses. In particular, our model omitted consideration of reproduction in the study sites, because we assumed most larvae arriving at the site were spawned elsewhere. As we obtain better, finer-scale estimates of larval retention and larval dispersal pathways (e.g., Saenz-Agudelo et al. 2011, Harrison et al. 2012), nearshore circulation (e.g., Drake et al. 2013, Harrison

et al. 2013), and planktonic larval mortality rates (White et al. 2014), it may be possible to include a realistic representation of closed population dynamics. This would likely have to be done by creating a spatial IPM that simultaneously models dynamics at multiple coastal sites that are linked by dispersal. Of course, in terrestrial systems or for marine species without a dispersive larval stage (e.g., Sanford and Worth 2009), there would be a much lower hurdle to completing the demographic loop in this type of model. One additional feature that would need to accompany any inclusion of reproduction is density-dependence in some vital rate. For the rockfish system, this would likely take the form of density-dependent post-settlement mortality of juveniles (Johnson 2006*a, b*, 2007) but there could be density-dependence in adult abundance or in reproductive output as well. Without density dependence, a linear IPM will eventually exhibit geometric growth; for a state-space model this would likely result in the process error term (or perhaps other parameters) compensating for the divergence between model and data, and any model projections beyond a few time steps would be suspect. Though our model was linear, recruitment entered as a subsidy term, so there was a deterministic equilibrium.

The state-space IPM approaches we have described here afford the opportunity for accurate model estimation from length–abundance time series data. This approach can strengthen ecological inference, allowing the estimation of unknown demographic parameters over appropriate spatial and temporal scales, and facilitating short-term predictions about transient system dynamics.

Acknowledgments

We thank L. Barnett for helpful discussions in early stages of this work and two anonymous reviewers for helpful comments that improved the manuscript. This work was supported by California Sea Grant (R/FISH-211) and the National Science Foundation (OCE-1435473, OCE-1260693) and is contribution number 436 from PISCO, the Partnership for Interdisciplinary Studies of Coastal Oceans, funded primarily by the David and Lucile Packard Foundation.

Literature Cited

Babcock, R. C., N. T. Shears, A. C. Alcala, N. S. Barrett, G. J. Edgar, K. D. Lafferty, T. R. McClanahan, and G. R. Russ. 2010. Decadal trends in marine reserves reveal differential rates of change in direct and indirect effects. Proceedings of the National Academy of Sciences USA 107:18256–18261.

Botsford, L. W., J. W. White, M. H. Carr, and J. E. Caselle. 2014. Marine protected area networks in California, USA. Pages 205–251 *in* M. L. Johnson and J. Sandell, editors. Advances in marine biology. Elsevier Academic Press, London, UK.

Briggs, J., K. Dabbs, M. Holm, J. Lubben, R. Rebarber, B. Tenhumberg, and D. Riser-Espinoza. 2010. Structured population dynamics: an introduction to integral modeling. Mathematics Magazine 83:243–257.

Brooks, S., A. Gelman, G. L. Jones, and X.-L. Meng, editors. 2011. Handbook of Markov chain Monte Carlo. Taylor and Francis, Boca Raton, Florida, USA.

Bruno, J. F., S. P. Ellner, I. Vu, K. Kim, and C. D. Harvell. 2011. Impacts of aspergillosis on sea fan coral demography: modeling a moving target. Ecological Monographs 81:123–139.

Burgess, S. C., K. J. Nickols, C. D. Griesemer, L. A. K. Barnett, A. G. Dedrick, E. V. Satterthwaite, L. Yamane, S. G. Morgan, J. W. White, and L. W. Botsford. 2014. Beyond connectivity: how empirical methods can quantify population persistence to improve marine protected-area design. Ecological Applications 24:257–270.

Carr, M. H., and D. C. Reed. 2015. Shallow rocky reefs and kelp forests. Pages 311–336 *in* H. Mooney and E. Zavaleta, editors. Ecosystems of California. University of California Press, Berkeley, California, USA.

Carr, M., and C. Syms. 2006. Recruitment. Pages 411–427 *in* L. G. Allen, D. J. Pondella II, and M. H. Horn, editors. The ecology of marine fishes: California and adjacent waters. University of California Press, Berkeley, California, USA.

Carr, M. H., C. B. Woodson, O. M. Cheriton, D. Malone, M. A. McManus, and P. T. Raimondi. 2011. Knowledge through partnerships: integrating marine protected area monitoring and ocean observing systems. Frontiers in Ecology and the Environment 9:342–350.

Casella, G., and E. George. 1992. Explaining the Gibbs sampler. American Statistician 46:167–174.

Caselle, J. E., B. P. Kinlan, and R. R. Warner. 2010*a*. Temporal and spatial scales of influence on nearshore fish settlement in the Southern California Bight. Bulletin of Marine Science 86:355–385.

Caselle, J. E., J. R. Wilson, M. H. Carr, D. P. Malone, and D. E. Wendt. 2010*b*. Can we predict interannual and regional variation in delivery of pelagic juveniles to nearshore populations of rockfishes (genus *Sebastes*) using simple proxies of ocean conditions? Reports of California Cooperative Oceanic Fisheries Investigations 51:91–105.

Caselle, J. E., S. L. Hamilton, D. M. Schroeder, M. S. Love, J. D. Standish, J. A. Rosales-Casián, and O. Sosa-Nishizaki. 2011. Geographic variation in density, demography, and life history traits of a harvested, sex-changing, temperate reef fish. Canadian Journal of Fisheries and Aquatic Sciences 68:288–303.

Castrejón, M., and A. Charles. 2013. Improving fisheries co-management through ecosystem-based spatial management: the Galapagos Marine Reserve. Marine Policy 38:235–245.

Caswell, H. 2001. Matrix population models: construction, analysis, and interpretation. Second edition. Sinauer Associates , Sunderland, Massachusetts, USA.

Chesson, P. 2012. Scale transition theory: its aims, motivations, and predictions. Ecological Complexity 10:52–68.

Chib, S., and E. Greenberg. 1995. Understanding the Metropolis–Hastings algorithm. American Statistician 49:22.

Clark, J. S. 2003. Uncertainty and variability in demography and population growth: a hierarchical approach. Ecology 84:1370–1381.

Clark, J. S., D. M. Bell, M. H. Hersh, M. C. Kwit, E. Moran, C. Salk, A. Stine, D. Valle, and K. Zhu. 2011. Individual-scale variation, species-scale differences: inference needed to understand diversity. Ecology Letters 14:1273–1287.

Coulson, T. 2012. Integral projections models, their construction and use in posing hypotheses in ecology. Oikos 121:1337–1350.

Crouse, D. T., L. B. Crowder, and H. Caswell. 1987. A stage-based population model for loggerhead sea turtles and implications for conservation. Ecology 68:1412–1423.

Cvitanovic, C., et al. 2013. Critical research needs for managing coral reef marine protected areas: perspectives of academics and managers. Journal of Environmental Management 114: 84–91.

Dennis, B., J. M. Ponciano, S. R. Lele, M. L. Taper, and D. F. Staples. 2006. Estimating density dependence, process noise, and observation error. Ecological Monographs 76:323–341.

de Valpine, P., and A. Hastings. 2002. Fitting population models incorporating process noise and observation error. Ecological Monographs 72:57–76.

Doak, D., P. Kareiva, and B. Klepetka. 1994. Modeling population viability for the desert tortoise in the western Mojave Desert. Ecological Applications 4:446–460.

Drake, P. T., C. A. Edwards, S. G. Morgan, and E. P. Dever. 2013. Influence of larval behavior on transport and population connectivity in a realistic simulation of the California Current System. Journal of Marine Research 71:317–350.

Easterling, M. R., S. P. Ellner, and P. M. Dixon. 2000. Size-specific sensitivity: applying a new structured population model. Ecology 81:694–708.

Eero, M., M. Vinther, H. Haslob, B. Huwer, M. Casini, M. Storr-Paulsen, and F. W. Köster. 2012. Spatial management of marine resources can enhance the recovery of predators and avoid local depletion of forage fish. Conservation Letters 5:486–492.

Ellner, S. P., and M. Rees. 2006*a*. Integral projection models for species with complex demography. American Naturalist 167:410–428.

Ellner, S. P., and M. Rees. 2006*b*. Stochastic stable population growth in integral projection models: theory and application. Journal of Mathematical Biology 54:227–256.

Freiwald, J. 2012. Movement of adult temperate reef fishes off the west coast of North America. Canadian Journal of Fisheries and Aquatic Sciences 69:1362–1374.

Froese, R., and D. Pauly. 2012. FishBase. http://fishbase.org

Gelfand, A. E., S. Ghosh, and J. S. Clark. 2013. Scaling integral projection models for analyzing size demography. Statistical Science 28:641–658.

Gelman, A., and K. Shirley. 2011. Inference from simulations and monitoring convergence. Pages 163–174 *in* S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors. Handbook of Markov chain Monte Carlo. Taylor and Francis, Boca Raton, Florida, USA.

Ghosh, S., A. E. Gelfand, and J. S. Clark. 2012. Inference for size demography from point pattern data using integral projection models. Journal of Agricultural, Biological and Environmental Statistics 17:641–677.

González, E. J., and C. Martorell. 2013. Reconstructing shifts in vital rates driven by long-term environmental change: a new demographic method based on readily available data. Ecology and Evolution 3:2273–2284.

González, E. J., C. Martorell, and B. M. Bolker. 2016. Inverse estimation of integral projection model parameters using time series of population-level data. Methods in Ecology and Evolution 7:147–156.

Gordon, N. J., D. J. Salmond, and A. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings F (Radar and Signal Processing) 140:107–113.

Green, P. J., and A. Mira. 2001. Delayed rejection in reversible jump Metropolis–Hastings. Biometrika 88:1035–1053.

Green, K. M., A. P. Greenley, and R. M. Starr. 2014. Movements of blue rockfish (*Sebastes mystinus*) off central California with comparisons to similar species. PLoS ONE 9:e98976.

Hallacher, L. E., and D. A. Roberts. 1985. Differential utilization of space and food by the inshore rockfishes (Scorpaenidae: Sebastes) of Carmel Bay, California. Environmental Biology of Fishes 12:91–110.

Harrison, H. B., et al. 2012. Larval export from marine reserves and the recruitment benefit for fish and fisheries. Current Biology 22:1023–1028.

Harrison, C. S., D. A. Siegel, and S. Mitarai. 2013. Filamentation and eddy–eddy interactions in marine larval accumulation and transport. Marine Ecology Progress Series 472:27–44.

Hutchings, J. A., and D. J. Fraser. 2008. The nature of fisheries- and farming-induced evolution. Molecular Ecology 17: 294–313.

Johnson, D. W. 2006*a*. Density dependence in marine fish populations revealed at small and large spatial scales. Ecology 87:319–325.

Johnson, D. W. 2006*b*. Predation, habitat complexity, and variation in density-dependent mortality of temperate reef fishes. Ecology 87:1179–1188.

Johnson, D. W. 2007. Habitat complexity modifies post-settlement mortality and recruitment dynamics of a marine fish. Ecology 88:1716–1725.

Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. Journal of Basic Engineering 82:35–45.

Key, M., A. D. MacCall, T. Bishop, and B. Leos. 2005. Stock assessment of the Gopher Rockfish (*Sebastes carnatus*). California Department of Fish and Game, Santa Cruz, CA, USA.

Key, M., A. D. MacCall, J. Field, D. Aseltine-Neilson, and K. Lynn. 2008. The 2007 assessment of blue rockfish (*Sebastes mystinus*) in California. California Department of Fish and Game, Santa Cruz, CA, USA.

Kirlin, J., M. Caldwell, M. Gleason, M. Weber, J. Ugoretz, E. Fox, and M. Miller-Henson. 2013. California's Marine Life Protection Act initiative: supporting implementation of legislation establishing a statewide network of marine protected areas. Ocean and Coastal Management 74:3–13.

Knape, J., and P. de Valpine. 2012. Fitting complex population models by combining particle filters with Markov chain Monte Carlo. Ecology 93:256–263.

Leet, W. S., C. M. Dewees, R. Klingbeil, and E. J. Larson, editors. 2002. California's living marine resources: a status report. California Department of Fish and Game Resources Agency, Sacramento, CA, USA.

Love, M. S., M. Yoklavich, and L. Thorsteinson. 2002. The rockfishes of the northeast Pacific. University of California Press, Berkeley, California, USA.

Madin, J. S., T. P. Hughes, and S. R. Connolly. 2012. Calcification, storm damage and population resilience of tabular corals under climate change. PLoS ONE 7:e46637.

Meinhold, R. J., and N. D. Singpurwalla. 1983. Understanding the Kalman filter. American Statistician 37:123–127.

Merow, C., J. P. Dahlgren, C. J. E. Metcalf, D. Z. Childs, M. E. K. Evans, E. Jongejans, S. Record, M. Rees, R. Salguero-Gómez, and S. M. McMahon. 2014. Advancing population ecology with integral projection models: a practical guide. Methods in Ecology and Evolution 5:99–110.

Metcalf, C. J. E., S. M. McMahon, R. Salguero-Gómez, and E. Jongejans. 2012. IPMpack: an R package for integral projection models. Methods in Ecology and Evolution 4:195–200.

Methot Jr., R. D., and C. R. Wetzel. 2013. Stock synthesis: a biological and statistical framework for fish stock assessment and fishery management. Fisheries Research 142:86–99.

Miller, D. J., and J. J. Geibel. 1973. Summary of blue rockfish and lingcod life histories; a reef ecology study; and giant kelp, *Macrocystis pyrifera*, experiments in Monterey Bay, California. California Department of Fish and Game Resources Agency, Sacramento, CA, USA.

Moffitt, E. A., J. W. White, and L. W. Botsford. 2013. Accurate assessment of marine protected area success depends on metric and spatiotemporal scale of monitoring. Marine Ecology Progress Series 489:17–28.

Rees, M., and S. P. Ellner. 2009. Integral projection models for populations in temporally varying environments. Ecological Monographs 79:575–594.

Rees, M., D. Z. Childs, and S. P. Ellner. 2014. Building integral projection models: a user's guide. Journal of Animal Ecology 83:528–545.

Saenz-Agudelo, P., G. P. Jones, S. R. Thorrold, and S. Planes. 2011. Connectivity dominates larval replenishment in a coastal reef fish metapopulation. Proceedings of the Royal Society B 278:2954–2961.

Sanford, E., and D. J. Worth. 2009. Genetic differences among populations of a marine snail drive geographic variation in predation. Ecology 90:3108–3118.

Scholz, A., K. Bonzon, R. Fujita, N. Benjamin, N. Woodling, P. Black, and C. Steinback. 2004. Participatory socioeconomic analysis: drawing on fishermen's knowledge for marine protected area planning in California. Marine Policy 28:335–349.

Starr, R. M., J. M. Cope, and L. A. Kerr. 2002. Trends in fisheries and fishery resources associated with the Monterey Bay National Marine Sanctuary from 1981–2000. California Sea Grant College Program, University of California, San Diego, California, USA.

Starr, R. M., M. Carr, D. Malone, A. Greenley, and S. McMillan. 2010. Complementary sampling methods to inform ecosystem-based management of nearshore fisheries. Marine and Coastal Fisheries 2:159–179.

Starr, R. M., et al. 2015. Variation in responses of fishes across multiple reserves within a network of marine protected areas in temperate waters. PLoS ONE 10:e0118502.

Walters, C. 1986. Adaptive management of renewable resources. McMillan, New York, New York, USA.

White, J. W., and J. E. Caselle. 2008. Scale-dependent changes in the importance of larval supply and habitat to abundance of a reef fish. Ecology 89:1323–1333.

White, J. W., L. W. Botsford, and E. A. Moffitt. 2010. Decision analysis for designing marine protected areas for multiple species with uncertain fishery status. Ecological Applications 20:1523–1541.

White, J. W., L. W. Botsford, M. L. Baskett, L. A. Barnett, R. J. Barr, and A. Hastings. 2011. Linking models with monitoring data for assessing performance of no-take marine reserves. Frontiers in Ecology and the Environment 9:390–399.

White, J. W., L. W. Botsford, A. Hastings, M. L. Baskett, D. M. Kaplan, and L. A. K. Barnett. 2013. Transient responses of fished populations to marine reserve establishment. Conservation Letters 6:180–191.

White, J. W., S. G. Morgan, and J. L. Fisher. 2014. Planktonic larval mortality rates are lower than widely expected. Ecology 95:3344–3353.

Wilcox, C., and C. Pomeroy. 2003. Do commercial fishers aggregate around marine reserves? Evidence from Big Creek Marine Ecological Reserves, Central California. North American Journal of Fisheries Management 23:241–250.

Yau, A. J., H. S. Lenihan, and B. E. Kendall. 2014. Fishery management priorities vary with self-recruitment in sedentary marine populations. Ecological Applications 24:1490–1504.

Zuidema, P. A., E. Jongejans, P. D. Chien, H. J. During, and F. Schieving. 2010. Integral Projection Models for trees: a new parameterization method and a validation of model output. Journal of Ecology 98:345–355.

Supporting Information

Additional Supporting Information may be found online at: http://onlinelibrary.wiley.com/doi/10.1002/eap.1398/suppinfo

Data Availability

All model code and example data are available online: https://github.com/jwilsonwhite/IPM_statespace