A KALMAN FILTER FOR ENVIRONMENTAL RISK: SPATIO-TEMPORAL VARIATION IN SEA TURTLE BYCATCH RATES

ABSTRACT

We extend a Kalman filter model for count data to reflect observations irregularly spaced in time. The model is applied to estimate risk of rare events; bycatch of endangered leatherback turtles in a swordfish fishery off Oregon and California. Incidental takes of turtles are modeled as a Poisson process. The results suggest that take risk is better estimated by regarding the entire history of observations, contrary to only considering data from the previous season as in the current management regime. The conservatory measures currently enforced in the fishery may be overly stringent.

JEL Classification: Q22, Q28, Q57, Q58.

Keywords: Kalman filter, Poisson process, government conservation policy, endangered species, biodiversity conservation.

INTRODUCTION

The drift gillnet fishery on swordfish (*Xiphias gladius*) along the southwestern U.S. coast is currently managed through a seasonal closure of a substantial part of the fishing grounds. The closure area was established in order to protect leatherback turtles (*Dermochelys coriacea*) from incidental capture. The statistical analysis that informs the management decisions in the fishery implicitly assume data on incidental turtle takes from the most recent observations provide the best information regarding current incidental capture risk. Efforts to predict unobserved turtle takes do not build upon a probability model; instead linear extrapolation is used to predict unobserved takes from observed takes. In other words, full correlation between the observed and unobserved take rate is implicitly assumed. We estimate the risk of incidental takes of sea turtles and find, contrary to the principles inherent in the current management regime, that the risk is best understood by considering looking at the entire recorded history of effort and bycatch. We argue that incidental takes of sea turtles have an element of randomness in them, and that management decisions may be improved when based on assessments of distributions of possible outcomes instead of point estimates based on limited data, as in the current regime.

In 2001 a time-and-area closure was established for a large portion of the drift gillnet fishery in California and Oregon. Before the closure, roughly one third of the effort units (sets) occurred inside the now-closed area and most (roughly 80%) of that effort occurred during the closed period. In subsequent years there was substantial loss in average annual revenues in the fishery. On the other hand, dwindling stocks of Pacific leatherback turtles have raised concerns about the continued existence of the species. Both of these issues need to be addressed through the management of the fishery.

The U.S. Federal Endangered Species Act requires that each federal agency “shall ensure that any action authorized, funded, or carried out by such agency is not likely to jeopardize the continued existence of any endangered or threatened species or result in the destruction or adverse modification of critical habitat of such species” (Section 7 (a) (2)). In order to comply with the Endangered Species Act, the responsible federal agencies have to analyze the available data on leatherback turtle incidental takes (or ‘interactions’ in the official language) and the level and distribution of effort in the fishery. This analysis informs the decisions on how to manage the incidental take risk posed by the fishing activity. Data from an observer program mandated by the 1988 amendments to the Marine Mammal Protection Act are used in this analysis. Almost 20 percent of the effort has been observed through the observer program. A key question regarding the observer data is the degree of independence between turtle takes in the observed part of the effort versus turtle takes in the unobserved effort. Implicit in the present approach to understand incidental takes is a belief that the process governing fluctuations in takes is a case of pure uncertainty, implying that no probability model can be used to explain leatherback takes. Further, only the current season's observations are used to predict takes in the unobserved part of the effort, and the takes in the unobserved part of the effort are predicted assuming perfect linear correlation with the number of leatherback takes in the observed effort; naïve extrapolation predicts unobserved takes.

We wish to suggest improvements to the statistical analysis of the available data. We model the rare event of an incidental turtle take as a Poisson process and estimate the risk of takes by using and extending a Kalman filter for count data models. Risk may differ across time and space. In our analysis observations from previous seasons may be used to predict unobserved turtle takes, predictions which are best described with probability distributions rather than point estimates. Our results suggest that predictions used to justify the 2001 closure may overstate the leatherback incidental take risk, which potentially leads to excessively stringent conservation measures.

The knowledge and understanding of the leatherback in general and its stock dynamic and migration pattern is limited; no analysis can fully gauge the impact of management actions on the leatherback population. There may be reason to believe that restrictions on U.S. catches of swordfish displace catches to non-U.S. fisheries where bycatch rates of leatherbacks are higher; Rausser et al. (2008) provides an example of this transfer effect from the Hawaiian longline fishery. The effort to protect leatherback turtles in the U.S. may lead to a larger short term global level risk of extinction. (By contrast, long term effects may include protection in other jurisdictions partly induced by protection in U.S. waters.)

We wish to extend the analysis and take a look at the economics implications of our results. How big a share of the effort in the fishery needs to be observed in order to have a certain level of confidence in the predicted total level of turtle takes, and would the costs of such levels of data collection allow for an economically viable fishery? Are we able to infer necessary conditions on the efficiency of the fishery given a certain level of observation costs? And ultimately, can we unite the economic and conservational requirements at any level of activity in the fishery? More work is required to address these questions.

DATA AND THE OBSERVER PROGRAM

The U.S. National Marine Fisheries Service established an observer program as mandated by the 1988 amendments to the Marine Mammal Protection Act. Since 1990, drift gillnet fishing effort has been observed in the U.S. exclusive economic zone from the waters off San Diego, California to the waters off Oregon. Observers record catch and bycatch by taxon for fish, marine mammals, and sea turtles, collect specimens, and record data on environmental conditions and over 10 different net characteristics (National Marine Fisheries Service 1997). Geographical coordinates are also recorded.

Vessels are selected on an opportunistic basis, with the hope that the observed sample will be representative of all the effort that occurred in a given year. Vessels are required to carry an observer about 20 percent of the time on a rotating basis; thus a vessel which just took an observer would not be required to carry another observer until it would approach their 20 percent requirement. Observer coverage averaged approximately 16 percent from 1991-99.

The sample of observed effort is not a true simple random sample of overall effort for the drift gillnet fishery in a given season, as the selection of vessels to be observed is by nonrandom choice, and the observed sets of effort are spatially and temporally clustered according to which vessels' trips are selected. Nonetheless, it is conceivable that given the rare and unpredictable nature of leatherback turtle bycatch, the pattern of takes which occurs over time in the observed portion of drift gillnet effort is highly similar to what would be seen in a true random sample. It is useful to conduct the thought experiment regarding the relationship between leatherback takes in the observed and unobserved portions of effort if the observed share of effort were chosen as a simple random sample of overall effort. In this case, after conditioning on factors which affect (the mean) leatherback take risk for the fishery as a whole, the number of takes in the observed and unobserved portions of effort would be independent by design of the sampling method.

Departure from full independence could arise if unobserved fishermen faced a systematically different risk of leatherback takes, or if the variance of within-trip leatherback take risk was smaller than the variance of leatherback take risk between trips. The second potential departure from independence could, in principle, be tested empirically by comparing between-trip variation in leatherback takes to within-trip variation. The first potential departure from independence is more problematic to assess, as unobserved effort is, by its very nature, not available for empirical investigation. On the grounds that drift gillnet fishermen have little to gain and much to lose from catching leatherback turtles, we see no prior reason to assume an elevated take risk in the unobserved portion of effort.[1] We adopt a working hypothesis for the balance of this paper that the observer sample may be reasonably treated as observationally equivalent to a simple random sample of all effort for the season, recognizing that whether this is the case is an open question.

The data consist of 7733 observed sets over 1331 different fishing trips. The number of sets per trip range from 1 to 19. The observations were made during the seasons 1990 – 2006; seasons mainly stretch from August through January, although fishing began as early as May in some years. The observations contain the number of leatherback interactions, date and geographical coordinates for each set. Over the 7733 sets, 23 interactions were observed. That is, the mean of the interactions is .0030 while the standard deviation is .0545. The temporary closure (Aug. 15 – Nov. 15) of parts of the fishing ground was enforced after the 2000 season, which ended on Jan. 31, 2001. Table 1 summarizes how the sets were distributed relative to the closure. The 'Pre Closure' column accounts for the seasons 1990 – 2000; 11 out of 17 seasons (64.7%). 77.9% of the sets were fished before the closure; the activity in the fishery has roughly gone down 50% after the closure. The average number of sets per season was 454 over all seasons, 547 before, and 284 after the closure. The share of sets fished inside the closed area went down substantially after the closure, while the share of sets fished in the closure period only went down slightly. Evidently, most of the 'in-area' sets (2145) were fished 'in-season' (1820). Note, however, that almost no sets were fished in the area but out of season after the closure. There are two possible explanations. Either the fishery is only profitable in the closure area between Aug. 15 and Nov. 15, or the fishermen prefer to fish close to their home port and those located in ports well within the area left the fishery after the closure. (84.8% [the ratio of sets fished in area during the season versus all sets fished in area pre closure.] of their 'preferred' season was closed.) Table A1 in the Appendix breaks the data into seasons; Table A2 shows how the turtle takes were distributed across seasons.

[1] There are potential scenarios where Leatherback take risk could be systematically higher in the unobserved portion of the effort. For example, if swordfish catch per unit of effort and Leatherback bycatch per unit of effort exhibited a high degree of spatial-temporal correlation, it is possible that fishermen carrying an observer would accept a lower swordfish catch per unit of effort to reduce the risk of observed Leatherback bycatch. Due to the extremely rare incident of Leatherback bycacth, we are skeptical that fishermen can assess Leatherback take risk with a sufficient degree of precision for this scenario to be viable.

Table 1: Data summary

	Tot. Obs.	Pre Closure	Post Closure
Tot. observations	7733	6024 (77.9%)	1709 (22.1%)
In Area	2192 (28.3%)	2145 (35.6%)	47 (2.75%)
In Season	4378 (56.6%)	2482 (57.8%)	896 (52.4%)
In Area & In Season	1826 (23.6%)	1820 (30.2%)	6 (0.351%)

THE KALMAN FILTER AND ITS USES

The standard model for the Kalman Filter is called the state-space model and consists of the observation equation, the state equation, and assumptions on the variance of error terms in both equations. The observation equation models the relationship between the observed variable and the unobserved state of the world. The relationship could be any general function, but is often modeled as linear. The state equation models the relationship between the current state of the world and previous states. When we observe y_t and let μ_t be the state, we get the following equations.

$$y_t = F\mu_t + \varepsilon_t$$

$$\mu_{t+1} = f\mu_t + \eta_t$$

The first equation is the observation equation; F is a known matrix, ε_t is the normal, zero mean error term with variance H_t. The second equation is the state equation; f is a known matrix, and η_t is the normal, zero mean error term with variance Q_t. See, for example, Meinhold and Singpurwalla (1983; p. 123) or Harvey (1989; pp. 100-101) for further details on the state space model. When the model includes explanatory variables, these can be introduced by letting f be a so-called link function (Harvey 1989; p. 418). Compared to the standard linear regression model, μ_t is the counterpart to the regression coefficients. As the coefficients are the output from a linear regression, the state vector is the output from the Kalman filter (Anderson and Moore 1979; p. 41). It is obvious from the state equation that the filter lets the coefficients vary while they are fixed in the linear regression model. The filter consists of two sets of equations; prediction equations and updating equations. The prediction equations predict the state variables at the next observation and can usually be extended to predict any number of steps into the future. The predictions are updated to reflect data as they become available through the updating equations.

The recursive structure of the state-space model is important; one can imagine that estimations are carried out recursively. The derivation of the estimators may be based on either classical frequentist (Eubank 2006, Harvey 1989) or Bayesian principles (Meinhold and Singpurwalla 1983); which is a question of faith. In the classical formulation, the filter lets likelihood functions be calculated recursively without using large matrices (Jones 1993; p. 78). When estimators are established the question remains of 'initial conditions.' The initial conditions are imposed on the variables prior to the first observation. A natural question is what impact the conditions have on the estimates. Many texts are vague on this point (Eubank 2006), but at least it is possible to impose an 'improper,' 'uninformative' distribution on the variables prior to the first observation. Such distributions carry no information (every possible value has the same, positive probability), and are not proper distributions (they do not sum to one). The Kalman filter with an improper prior is often called a diffuse Kalman filter (Eubank 2006; p. 108). Initial conditions may also reflect knowledge or information about the state variables independent of the time series data used in the estimation. The state equation naturally gives rise to a one-step prediction problem, but the problem is still termed a filtering problem (Anderson and Moore 1979; p. 37). The Kalman filter yields a predictive distribution conditional on available data.

THE KALMAN FILTER MODEL OF LEATHERBACK INCIDENTAL TAKE RISK

The observer data cannot separately identify the stock density (of turtles) from catchability; hence we resort to a reduced form model. There is a base level risk μ_t of turtle take at set t. The risk varies, however, with the explanatory variables x_t; the risk of an incidental turtle take at set t is modeled as the product of the base level risk and the link function $\exp(x_t\delta)$ which can adjust to account for covariates which are believed to affect take risk.

$$\mu_t(x_t) = \exp(x_t\delta)\mu_t$$

$\delta = [\delta\ ,\ldots,\delta_n]'$ are parameters related to the covariates $x_t = [x_{1t},\ldots,x_{nt}]$. The base level risk μ_t is nonnegative. The exponential link function is always positive and the turtle take risk $\mu_t(x_t)$ is thus nonnegative. The base level risk is the unobservable state variable in the model.

Again, $\mu_t(x_t)$ is the risk of an incidental turtle take on set t. (Incidental takes of two or more turtles on one set may happen, but has not yet been observed; our model does however allow more than one turtle take per set.) We have the observation equation

$$y_t = z_t \mu_t(x_t) + \varepsilon_t$$

where z_t is the amount of effort (number of sets) in observation t. In the observer data, effort is observed on the set level and $z_t = 1$. Extension to data with multiple sets per observation is straightforward. We impose the conditional Poisson distribution on the turtle observations given the covariates;

$$(y_t \mid x_t) \sim Poisson(z_t \mu_t(x_t))$$

A conjugate prior for the Poisson distribution is the Gamma distribution (Harvey 1989; p. 351), that is, that the observations are Poisson distributed consists with a Gamma distributed base level risk. The predictive distribution for the base level risk on set t given observations on explanatory variables (X_{t-1}) and incidental takes (Y_{t-1}) through set $t-1$ is then denoted

$$(\mu_t \mid X_{t-1}, Y_{t-1}) \sim \Gamma(a_{t|t-1}, b_{t|t-1})$$

The Gamma distribution parameters $a_{t|t-1}$ and $b_{t|t-1}$ are estimated from the data in a maximum likelihood estimation.

Imposing the Poisson distribution on the observations may seem a strong assumption; the mean and variance of the Poisson distribution are identical. The mean and variance of turtle takes in the observer data are identical, however, and we find the Poisson distribution appropriate.

MAXIMUM LIKELIHOOD ESTIMATION

A maximum likelihood estimation procedure for the state space model with count data is developed by Harvey & Fernandes (1989, see also Harvey 1989; pp. 350-353, 418-420). The model treats the data as observed at regular time intervals. The drift gillnet fishery is, however, a seasonal fishery and time intervals between observations vary greatly; we extend the model to reflect the unequal spacing of the observations. Lambert (1996) extends the Harvey & Fernandes (1989) model for count data observed at unequally spaced times in a Bayesian framework. As far as we can tell, his approach is formally identical to ours. We maintain Harvey & Fernandes (1989) classical formulation, however, as it may be more familiar to some researchers. A key feature of the Kalman filter is that recent observations potentially should carry more weight in the estimators and predictors than earlier observations; when time passes earlier observations may loose their information value. The process may be called data discounting and involves the estimation of a data discount rate ω. Contrariwise, Lambert (1996; p. 37) comprehends ω as a 'serial association' coefficient. When forecasting the parameter values with data through observation t, the mean of the predicting distribution should be equal to the mean of the distribution estimated when observation t became available (adjusted for potential changes in the covariates). At the same time, the variance of the predicting distribution should increase with time if confidence is lost. While Lambert's (1996) formulation eloquently displays the mean-preserving-but-variance-increasing property in the predictive equations (p. 33), his updating equations (p. 34) are less immediate. Again, we prefer Harvey & Fernandes' (1989) classical formulation, where the updating equations (p. 408; p. 412) take on a more recognizable form and are readily interpreted.

Let s_t denote the date on which set t was fished and let $\omega \in (0,1]$ denote a rate of time discount. The parameters in the predictive Gamma distribution are defined recursively as follows:

$$\begin{aligned} a_{t|t-1} &= \omega^{s_t - s_{t-1}} a_{t-1} \\ b_{t|t-1} &= \omega^{s_t - s_{t-1}} b_{t-1} \exp(-x_t \delta) \end{aligned} \tag{1}$$

When set t is observed, the parameters are updated and form the basis for the predictive distribution for the next set. The parameters are updated as follows:

$$\begin{aligned} a_t &= \omega^{s_t - s_{t-1}} a_{t-1} + y_t = \sum_{j=1}^{t} \omega^{s_t - s_j} y_j = EWMA(y) \\ b_t &= \omega^{s_t - s_{t-1}} b_{t-1} + \exp(x_t \delta) = \sum_{j=1}^{t} \omega^{s_t - s_j} \exp(x_t \delta) = EWMA(\exp(x^T \delta)) \end{aligned} \tag{2}$$

where x^T is the transpose of x. The expressions for the updated parameters demonstrates that the parameters are exponentially-weighted moving averages of the entire history of observations on leatherback incidental takes (a_t) and risk adjusting exposures (b_t). Equations (1) and (2) are equivalent to the standard Kalman filter equations; see Harvey (1989; p. 163, p. 351) for details. Note that in the updating equations (2) it is the previous estimates a_{t-1} and b_{t-1} that are updated, and not the predicted values $a_{t|t-1}$ and $b_{t|t-1}$.

Harvey & Fernandes (1989) implicitly treats data as observed at regular time intervals, and the datum observed t periods ago is discounted with the factor ω^t in the exponentially-weighted moving averages; Interpretation of the discount factor then depends on the interval between observations. In equations (1) and (2), time enters explicitly and the discount factor is interpreted relative to how time is measured; we measure time in years and the discount factor is comprehended as a yearly discount rate. The exponentially-weighted moving average formula which describes the parameters of the posterior distribution of Leatherback take risk lends itself to a useful interpretation of the discount parameter. For $\omega = 1$, the mean of the posterior predictive distribution of the leatherback take risk is given by

$$E(\mu_t \mid Y_t) = a_t / b_t = \sum_{j=1}^{t} y_j \Big/ \sum_{j=1}^{t} \exp(x_j \delta)$$

which shows that the expected take risk conditional on observations through set t depends on an equally weighted averages of the entire history of observations; each observation carries equal weight. Contrary, for $\omega \to 0^+$, the mean of the posterior distribution of leatherback take risk is

$$\lim_{\omega \to 0^+} E(\mu_t \mid Y_t) = \lim_{\omega \to 0^+} a_t \Big/ \lim_{\omega \to 0^+} b_t = y_t / \exp(x_t \delta)$$

which is closer to the present approach of estimating the present year's leatherback take rate in the unobserved sets. The present approach sets the current year's unobserved take rate equal to the current year's observed take rate, ignoring earlier observations.

The Kalman filter lends itself to maximum likelihood estimation as the likelihood function is usually available in computable form; Harvey & Fernandes (1989; p. 409) develops the log-likelihood function for the filtering problem using the exponentially-weighted moving averages of the observations;

$$\begin{aligned}\ln L(\omega, \delta) = \sum_{t=\tau+1}^{T} \Big(& \ln \Gamma(a_{t|t-1} + y_t) - \ln \Gamma(a_{t|t-1}) + a_{t|t-1} \ln b_{t|t-1} \\ & - (a_{t|t-1} + y_t) \ln(1 + b_{t|t-1}) \Big)\end{aligned}$$

where τ is the first set in the data with $y_t > 0$ and T is the last observed set. Note that the likelihood function is only defined for observations after the first non-zero observation. The parameters that go into the function carry all available information, however. The maximum of the likelihood function was found using the Nelder-Mead simplex algorithm as implemented in Matlab's *fminsearch* function.

Harvey & Fernandes (1989; p. 409) suggest the following initial conditions: $a_0 = 0$, $b_0 = 0$, which implies $\Pr(\mu_0 = 0) = 1$; the initial distribution has zero variance. On the contrary, Lambert (1996; p. 34) suggest diffuse initial conditions with infinite variance. Both

Harvey (1989; pp. 137-140) and Eubank (2006; pp. 107-136) discuss initialization issues with the Kalman filter.

ALTERNATIVE MODEL SPECIFICATIONS

The observer data provide a rich set of information. It is, however, necessary to limit the number of explanatory variables in the estimation; the explanatory power of the relatively small number of leatherback takes is limited. The spatiotemporal closure of the main part of the fishing ground suggests that leatherback take risk varies through the year and over space. This may be reflected in the estimation through binary explanatory variables that relate to the seasonal closure and the closure area; x_{t1} indicates whether set t was fished inside the closure area and x_{t2} indicates whether the set was fished during the closed season (August 15 – November 15). x_{t3} indicates whether the set was fished inside the area during the closed season, *i.e.*, interaction between the area and season variable: $x_{t3} = x_{t1} \cdot x_{t2}$. Obviously, x_{t3} should be zero after the closure was implemented in 2001. The parameter vector $\delta = [\delta_1, \delta_2, \delta_3]'$ is associated to the explanatory variables and the subscripts of the parameters indicate which variable it relates to.

We estimate three specifications:

1. Unrestricted model;
2. Restricted model only including the area variable ($x_{t1}; \delta_2 = \delta_3 = 0$);
3. Restricted model without explanatory variables ($\delta = 0$).

Harvey & Fernandes (1989; p. 413) suggests checking specifications by computing the variance of the Pearson residual. The Pearson residual is obtained by applying a standard unit transformation to the data y_t, where the standardization uses the mean and standard deviation of the predictive distribution given time $t-1$ information:

$$v_t = \frac{y_t - E(y_t \mid \mu_{t-1})}{\sqrt{Var(y_t \mid \mu_{t-1})}}$$

A correctly specified model has a theoretical Pearson residual variance equal to 1; a sample Pearson residual variance close to 1 indicates a good fit.

MODELING THE ECONOMIC IMPACT OF REGULATION

If the model accurately describes turtle take risk with respect to time and area fished, fishers have no ability to influence turtle take risk besides timing (in or out of season) or location (in or out of the closure area). Further, if swordfish net revenue per set is R_0 while the closure area is open and R_1 under the closure and observer costs are C per set, the following questions can be addressed: What conditions on R_0, R_1 and C are necessary for it to be economically viable to open the conservation area with a mandatory level q of observer coverage? Are there ranges of values of R_0, C, and q for which reopening the area to fishing is viable from both economic and conservation standpoints?

ESTIMATION RESULTS

Estimation results are given in Table 2. The unrestricted model necessarily has the highest likelihood. To restrict the model to include no explanatory variables ($\delta = 0$) produces a likelihood ratio statistic that rejects the hypothesis that the explanatory variables can be excluded from the estimation (*p*-value is less than 1%). On the contrary, the likelihood ratio test for the model with only the area variable ($\delta_2 = 0, \delta_3 = 0$) cannot reject the hypothesis that the season and interaction variables can be excluded (*p*-value $\approx 30\%$).

While the likelihood ratio test results are inconclusive, assessing the Pearson residual variance suggests that the unrestricted model has the best fit. This implies that the restricted models exhibit *under-dispersion* relative to the variance predicted by the fitted time series model. ($Var(v) < 0.95$ suggests under-dispersion when the number of observations is 'high.') Failing to take account of the variation in leatherback take risk due to the location of effort may lead to an overestimation of the effect of the recent take rate on the current level of the take risk.

Both the unrestricted model and the model restricted to only include the area variable yield an estimated value of the time discount rate of 1. The take risk is thus best perceived taking the entire history of observations into account in estimation and forecasting. Since the covariates x_{ti} are dummy variables, the estimated effect from the different covariates are given by $\exp(\delta_i)$. They enter multiplicatively in the link function; *e.g.*, the rate of incidental turtle takes was approximately 14 times higher in the closure area between Aug. 15 and Nov.

15 than outside the area between Nov. 15 and Aug. 15 (δ_3 only captures the added effect from both x_{t1} and x_{t2} being 1 simultaneously).

Table 2: Comparison of alternative model specifications.

	Unrestricted	$\delta = 0$	$\delta_2 = 0, \delta_3 = 0$
Time Discount (ω)	1.0000	0.4580	1.0000
Closed Area (δ_1)	1.3856		2.4456
Closed Season (δ_2)	0.1752		
Interaction (δ_3)	1.1011		
ln L	-138.1979	-149.1102	-139.3577
$Var(v)$	0.9671	0.8736	0.9431
LR Statistic		21.8245	2.3196
Degrees of Freedom		3	2
p-value		0.0001	0.3136
$\exp(\delta_1)$	3.9971		11.5377
$\exp(\delta_2)$	1.1915		
$\exp(\delta_1 + \delta_2 + \delta_3)$	14.3236		

Figure 3 shows the estimated take rate of turtles per 1000 sets for the model without explanatory variables ($\delta = 0$). The take rate quickly settles down around 5, but varies quite a bit, between 1 and 8, until the closure area was established in 2001. There seems to be a clear break in the estimated take rate after the closing, and the rate converges towards zero after 2001. The estimated data discount factor is $\omega = 0.4580$ for this model, and the 'filtration' of earlier data is evident in Figure 3; particularly after the closure in 2001. A worst-case scenario with a take rate of 8 per 1000 sets and a high number of sets in a season, *e.g.*, 3,750 sets (757 sets were observed in 1993, representing approximately 20% of all sets) suggests that up to 30 turtles could be killed. This agrees rather well with the conclusions in the Biological Opinion (National Marine Fisheries Service 2000) which supported the seasonal closure. The estimated take rate in Figure 3 is, however, probably overstated due to omitted variable bias.

The establishment of the closure area may disturb our estimation; Figure 3 displays an apparent break in the take rate after the area was established. One way to deal with the problem is to control for the closure, as we do in the model with explanatory variables. Another possibility is to only regard data prior to the closure in the model without explanatory variables. When we estimate the $\delta = 0$ model with only the observations from seasons 1990 through 2000, we have $\omega = 0.9998$ ($\ln L = -147.3921; Var(v) = 0.9613$). The variance of the Pearson residuals suggests that the model fits the partial data set much better than the entire time series, where $Var(v) = 0.8736$. Data is discounted less with only the partial data; the partial-data estimate of the take rate reflects more of the regarded data compared to the estimate with the entire time series. Figure 4 shows the estimated take rate per 1000 sets with only the partial data set in the $\delta = 0$ model. The take rate settles around 4 takes per 1000 sets with relatively small variations compared to the variation in Figure 3. A worst-case scenario with 5 takes per 1000 sets in a 3,750-sets-season adds up to approximately 19 killed turtles, which is considerably less than what the Biological Opinion (National Marine Fisheries Service 2000) suggests. The graph in Figure 4 does not take all available information into account, however, and it does not reflect the belief that take risk varies through time and space. Let us turn to the best estimate; the full, unrestricted model applied to the full data set.

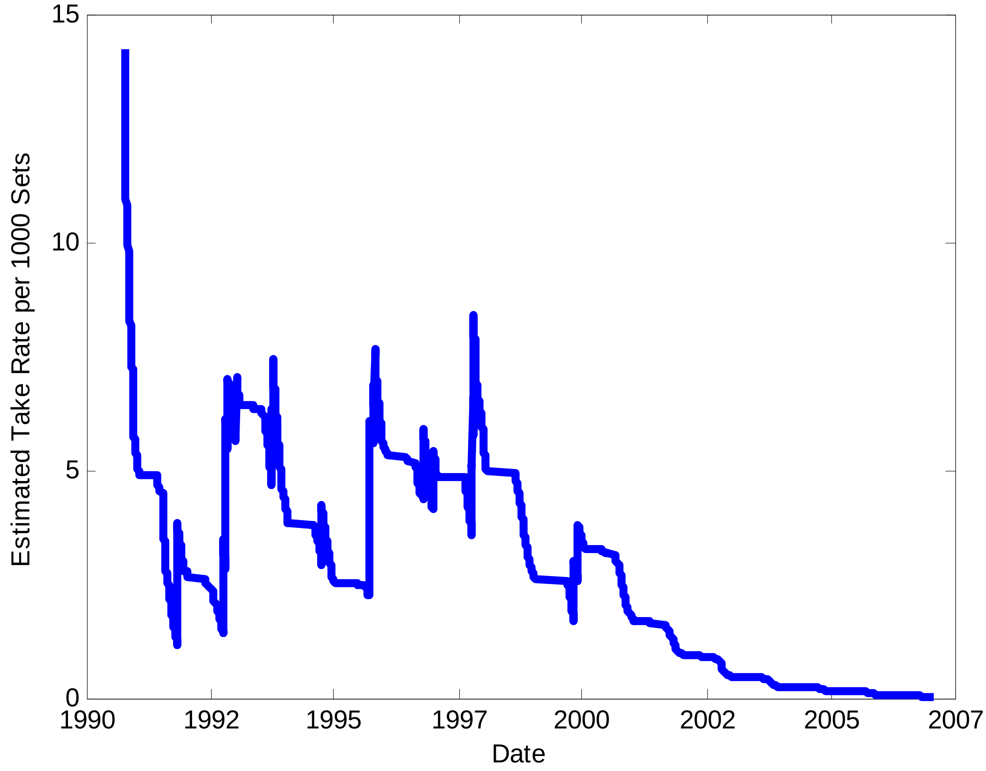

Figure 3: Estimated take rate per 1000 sets for the model without explanatory variables.

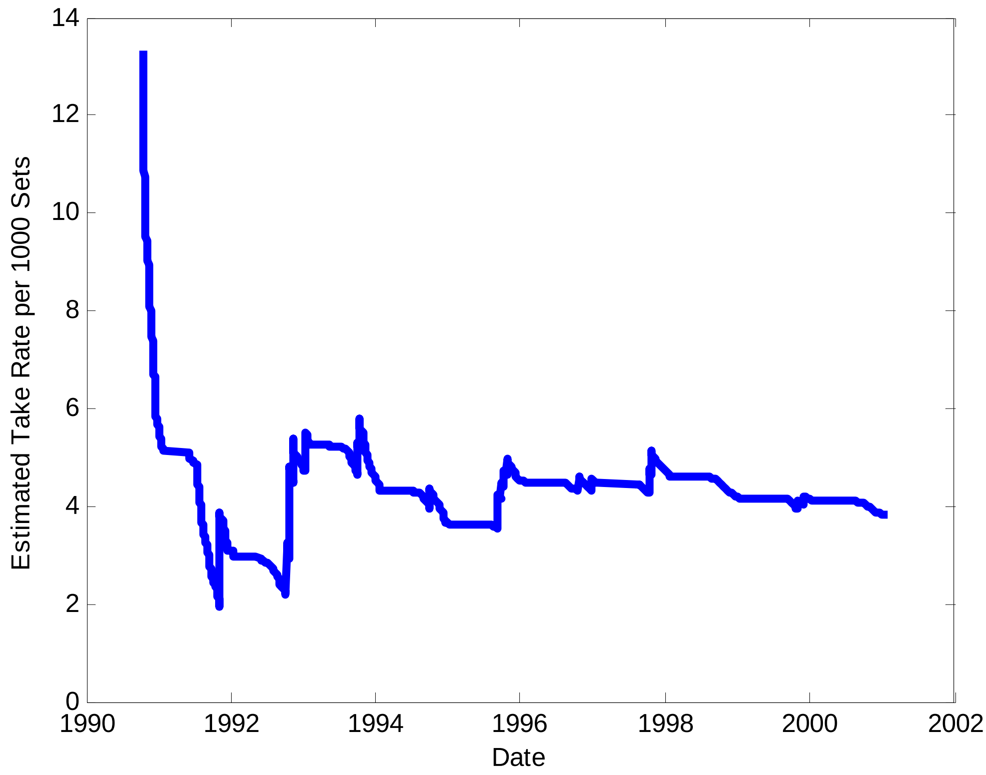

Figure 4: Estimated take rate per 1000 sets for the model without explanatory variables with partial data set (Seasons 1990-2000).

Figure 5 shows the estimated baseline turtle take rate per 1000 sets from the full, unrestricted model. The baseline rate corresponds to $\mu_t(x_t = 0) = \mu_t$; the take rate outside the turtle conservation area and outside the closure period. The take rate inside the area is $\exp(\delta_1) = 3.9971$ times higher than the baseline take rate; between Aug. 15 and Nov. 15 the take rate is $\exp(\delta_2) = 1.1915$ times higher; inside the area in the closure season the take rate is $\exp(\delta_1 + \delta_2 + \delta_3) = 14.3236$ times higher. (δ_3 only measures the increase in the take rate from sets both taking place in the area and in the period; these sets are already subjected to the increased take rate measured by δ_1 and δ_2; the total increase in the take rate with these sets relative to the baseline risk is measured by $\delta_1 + \delta_2 + \delta_3$.) *I.e.*, the current level of the baseline take rate is estimated to be 0.68; inside the area it is 2.7; between Aug. 15 and Nov. 15 it is 0.81; inside the area in the closure period it is 9.7. Combining these numbers with the 'average season' before the closure in 2001 suggests that approximately 10 turtles were killed in the average season. (See Table A1 for the pre closure average season figures.) If we take the effort-figures from the 1993 season as the worst-case scenario (again, see Table A1) with the estimated baseline take rate for the 1993 season, $\mu_{1993} = 0.67$ per 1000 sets, we get approximately 19 killed turtles. That is identical to the estimated 19 killed turtles in the worst-case scenario from the model without explanatory variables applied to the pre closure data set (1990-2000). (Note, however, that the take rate fell a bit during the 1993 season; the take rate is estimated at $\mu_{1992} = 0.88$ per 1000 sets prior to the 1993 season. Using 0.88 as the baseline take rate for the 1993 season results in approximately 24 turtle takes.)

Note the increased stability of the estimate in Figure 5 compared to the estimate in Figure 3. The stability is partly due to the estimated data discount factor $\omega = 1.0000$, which suggests that the entire history of effort should be equally weighted in the estimate of the take rate. However, some of the variation is now relegated to explain differences in the take rate across time and space.

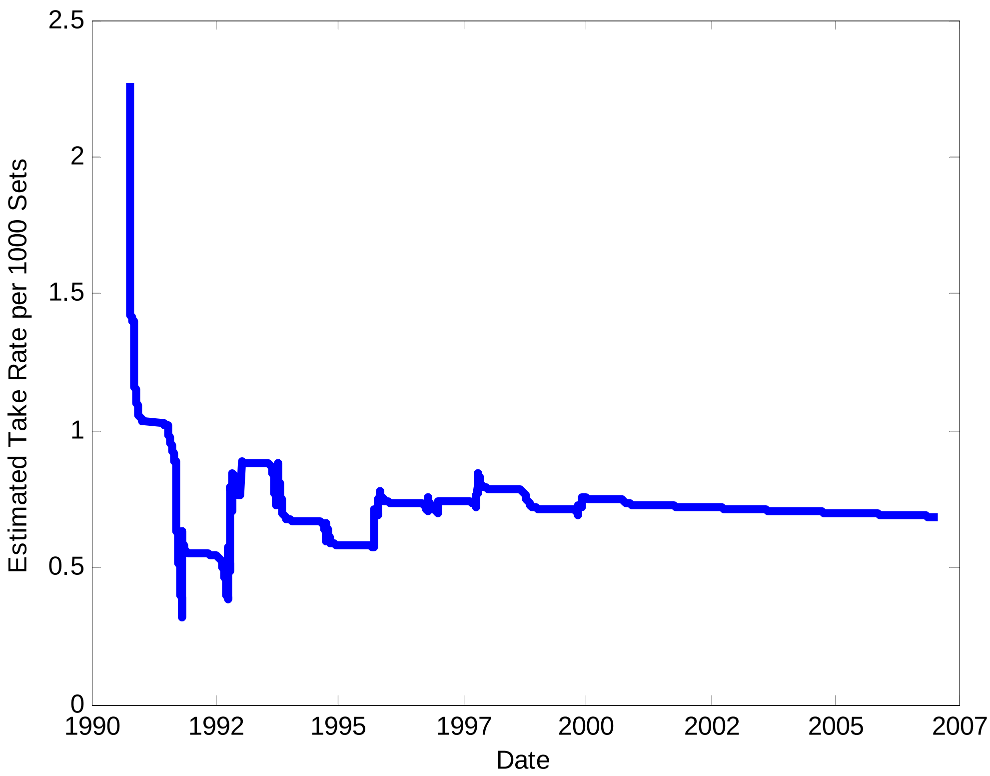

Figure 5: Estimated take rate per 1000 sets for the model with explanatory variables.

PREDICTION

Harvey & Fernandes (1989; p. 412) demonstrates that the predictive distribution of μ_t, given the data through observation y_{t-1}, is a negative binomial distribution with parameters $a_{t|t-1}$

and $b_{t|t-1}$. Figure 6 displays the predictive distribution for the observed number of takes in the 1995 season, where the parameters are our best estimates; that is, estimated from the full, unrestricted model. 1995 is an interesting season because it is the season with the highest number of turtle takes in the observer data (5 turtles was observed; see Table A2 in the appendix). As the graph shows, the probability of observing 5 turtle takes in the 1995 season was between 5 and 10%. The 5 observed turtles in 1995 was an unlikely event, but not necessarily rare.

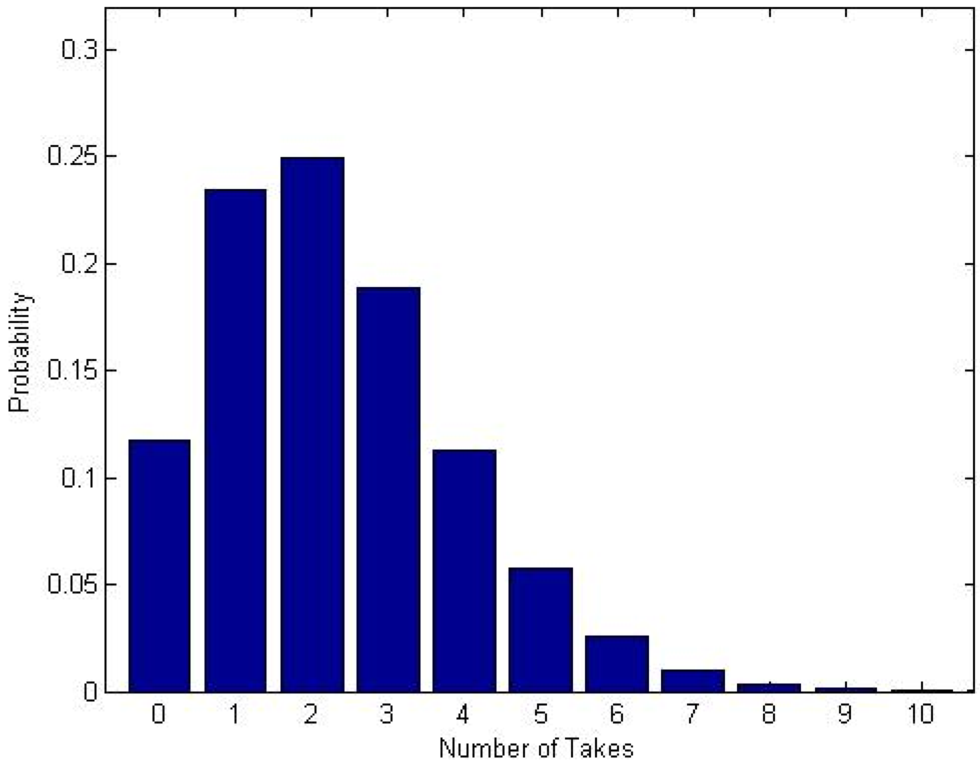

Figure 6: Predictive distribution for the number of observed 1995 leatherback takes.

Extrapolating the 5 observed turtle takes in the 1995 season, with 20% observer coverage and full covariance between the observed and unobserved effort, suggests that 25 turtles was taken in 1995. Figure 7 shows the predictive distribution for the total number of turtle takes, observed and unobserved, in the 1995 season which arises from our best estimate. The probability of totally 25 turtle takes is less than 2%; the probability of 25 turtle takes or more is less than 4%. Such events are rare, in other words. Moreover, our results suggest that the probability of maintaining such levels of turtle takes over several seasons is very small. The graph shows that it is much more likely that between 10 and 20 turtles were taken in 1995.

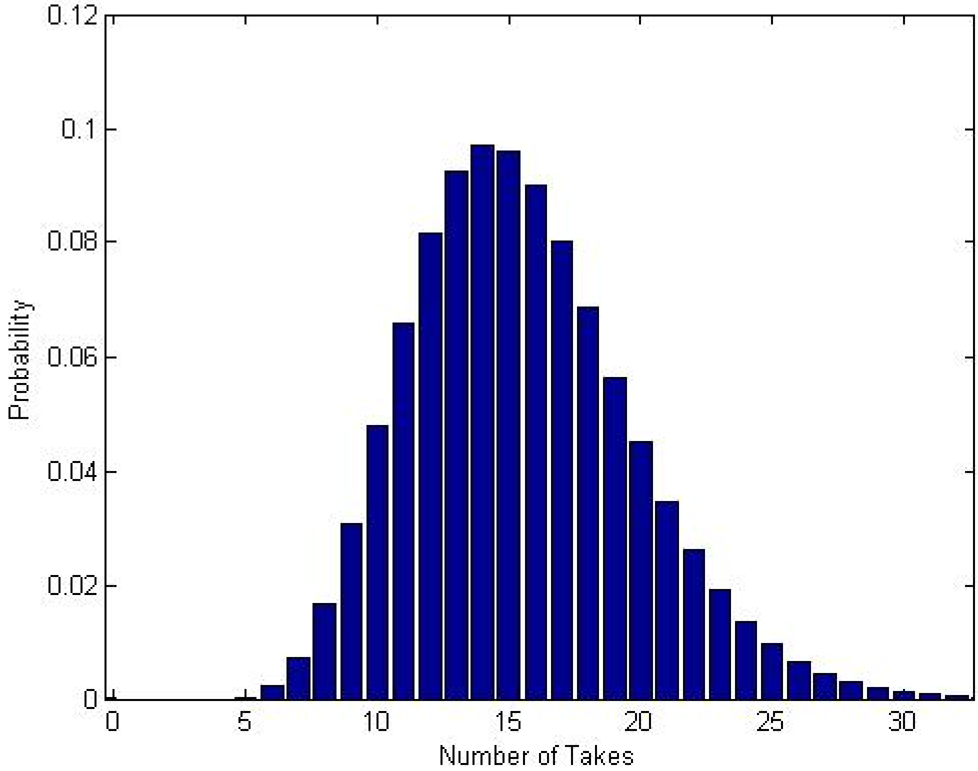

Figure 7: Predictive distribution for the total number of observed 1995 leatherback takes.

CONCLUSIONS

We used a Kalman filter approach to estimate the level of leatherback take risk over time based on observer data on fishing effort and leatherback takes from the drift gillnet fishery off California and Oregon. We represented take risk as a time-varying state variable that describes the rate parameter for a Poisson process, where the Poisson arrivals are incidental takes of leatherbacks. Our results confirm the earlier finding by Carretta et al. (2004) that historically there was significantly higher leatherback take risk in the area that was closed to the drift gillnet fleet fishing after 2001 during the period from Aug 15 through Nov 15.

The best fitting model appears to be one which includes explanatory variables for the time and area where fishing effort occurred, and which weights all past observations on drift gillnet leatherback takes and fishing effort equally. Our results suggest consideration of whether a longer time series of observations should be used for estimating the number of leatherback takes which occurred in the unobserved portion of each season’s effort. Any such predictions should, if possible, reflect variation in leatherback take risk over the times and areas where effort occurred.

APPENDIX

Table A1 breaks the set data into seasons. The number in parenthesis reflects the share relative to total observations in that season. The same is true for the mean and standard deviation rows; the numbers in the parenthesis in not the mean and standard deviation of the shares in the column above, but the share relative to the 'Total Observations' column.

Table A1: More Data Summary

Season	Tot. Obs.	In Area		In Season		Interaction	
1990	195	69	35.4 %	101	51.8 %	54	27.7 %
1991	477	168	35.2 %	257	53.9 %	162	34.0 %
1992	660	301	45.6 %	398	60.3 %	260	39.4 %
1993	757	416	55.0 %	390	51.5 %	342	45.2 %
1994	662	235	35.5 %	392	59.2 %	231	34.9 %
1995	587	221	37.6 %	356	60.6 %	179	30.5 %
1996	467	219	46.9 %	228	48.8 %	138	29.6 %
1997	748	253	33.8 %	453	60.6 %	212	28.3 %
1998	499	175	35.1 %	296	59.3 %	156	31.3 %
1999	528	48	9.1 %	350	66.3 %	48	9.1 %
2000	444	40	9.0 %	261	58.8 %	38	8.6 %
2001	323	5	1.5 %	201	62.2 %	5	1.5 %
2002	373	4	1.1 %	204	54.7 %	0	0.0 %
2003	295	0	0.0 %	176	59.7 %	0	0.0 %
2004	206	0	0.0 %	82	39.8 %	0	0.0 %
2005	228	23	10.1 %	93	40.8 %	0	0.0 %
2006	284	15	5.3 %	140	49.3 %	1	0.4 %
Mean	455	129	28.3 %	258	56.6 %	107	23.6 %
Std. Dev.	179	125	69.6 %	115	63.8 %	107	59.7 %
Pre Closure (1990-2000):							
Mean	548	195	35.6 %	317	57.8 %	165	30.2 %
Std. Dev.	154	108	70.3 %	96	62.1 %	90	58.8 %

Table A2 shows how the turtle takes were distributed over the seasons. All takes happened in the closure area during the closure period with 4 exceptions; one take (1992) happened outside the area and outside the period, two takes (1995 and 1999) happened outside the area but in the period, and one take (1996) happened inside the area but outside the period.

Table A2: Turtle Takes per Season

Season	Takes
1990	1
1991	1
1992	5
1993	2
1994	1
1995	5
1996	2
1997	4
1998	0
1999	2
2000	0
2001	0
2002	0
2003	0
2004	0
2005	0
2006	0

REFERENCES

Anderson, B.D.O. and J.B. Moore, 1979, Optimal Filtering, Dover Publications, Inc.

Carretta, J.V., T. Price, D. Peterson, and R. Read, 2004, Estimates of Marine Mammal, Sea Turtle, and Seabird Mortality in the California Drift Gillnet Fishery for Swordfish and Thresher Shark, 1996-2002, *Marine Fisheries Review*, 66(2), pp. 21-30.

Eubank, R.L., 2006, A Kalman Filter Primer, Chapman & Hall/CRC.

Harvey, Andrew C., 1989, Forecasting, structural time series models and the Kalman filter, Cambridge University Press.

Harvey, A.C. and C. Fernandes, 1989, Time series models for count or qualitative observations, *J. Bus, Statist.*, 7, pp. 407-423.

Jones, R.H., 1993, Longitudinal Data with Serial Correlation: A State-space Approach, Chapman & Hall

Lambert, P., 1996, Modelling of Repeated Series of Count Data Measured at Unequally Spaced Times, *Appl. Statist.*, 45(1), pp.31-38.

Meinhold, R.J. and N.D. Singpurwalla, 1983, Understanding the Kalman Filter, *Amer. Statist.*, 37(2), pp- 123-127.

National Marine Fisheries Service, 1997, Environmental Assessment of Final Rule to Implement the Pacific Offshore Cetacean Take Reduction Plan, under Section 118 of the Marine Mammal Protection Act, September 1997.

National Marine Fisheries Service, 2000, Biological Opinion Concerning Authorization to Take Listed Marine Mammals Incidental to Commercial Fishing Operations under Fishery Section 101(a)(5)(E) of the Marine Mammal Protection Act for the California/Oregon Drift Gillnet Fishery, October 23, 2000.

Rausser, Gordon C., Stephen F. Hamilton, Marty Kovach, and Ryan Stifter, 2008, Unintended Consequences: The Spillover Effects of Common Property Regulations, *CUDARE Working Papers* (University of California, Berkeley), 1053.