A comparison of gamma and lognormal maximum likelihood estimators in a sequential population analysis

Noel G. Cadigan and Ransom A. Myers

Abstract: We analyze the model used to assess most major commercial marine fish populations, namely, sequential population analysis (SPA). This model estimates population abundance by combining catch-at-age data with research surveys or commercial catch per unit effort indices of abundance. We examine two maximum likelihood estimators of SPA parameters. These estimators are based on assuming that the stock-size indices are from lognormal or gamma distributions. Using simulations, we find that both types of estimators can have significant biases; however, our results indicate that it is preferable to use the gamma model, because it tends to have lower bias and variability, even when the true distribution of the stock-size indices is lognormal.

Résumé : On trouvera ici un examen de l'analyse séquentielle de population (SPA), le modèle qui est utilisé pour évaluer la plupart des populations de poissons marins d'intérêt commercial. Ce type de modèle estime l'abondance de la population en combinant des données sur la capture à chacun des âges à des indices d'abondance basés sur des inventaires scientifiques ou sur le succès de la pêche commerciale par unité d'effort. Deux estimateurs de vraisemblance maximale des paramètres de la SPA ont été comparés; les deux estimateurs supposent que les indices de densité des stocks sont tirés de distributions lognormales ou gamma. Des simulations montrent que les deux types d'estimateurs peuvent introduire des erreurs systématiques importantes; cependant, d'après ces résultats, il est préférable d'utiliser le modèle gamma car il introduit moins de biais et de variation, même lorsque la véritable distribution des indices de densité des stocks est lognormale.

[Traduit par la Rédaction]

Introduction

Sequential population analysis (SPA) is an important model used to provide abundance information for many major commercial fisheries worldwide (Shepherd 1988). This model is so named because it sequentially combines commercial catches from all fish born in a given year, and uses research-survey estimates of relative abundance (indices) or commercial catch rates to estimate the present and past absolute fish stock abundance (numbers at age). SPA is the standard method used in the management of fish populations by both international organizations, e.g., ICES (International Council for the Exploration of the Sea) and NAFO (Northwest Atlantic Fisheries Organization), and national governments, e.g., Canada and the U.S.A. Versions of this method are also used to analyze populations of seals (Roff and Bowen 1983), whales (Cooke 1985), and terrestrial mammals (McCullough 1979). There are serious known biases in the method presently used that have led to overestimation of abundance and overexploitation of some populations (ICES 1991; Cook et al. 1991). Even so, SPA has received relatively little statistical investigation.

The standard method of estimating SPAs is to use nonlinear least squares with transformed survey indices. Transformations are used so that model errors appear homogeneously distributed; however, another reason is that nonlinear least squares is relatively easy to use. A log transformation in conjunction with a normal-error assumption is commonly used, which is equivalent to assuming that the untransformed survey indices have a lognormal distribution. However, with the exception of Patterson (1999), the appropriate distribution for modeling the survey indices has received little attention. Part of the reason for the lack of statistical attention is that many researchers examine model residuals; they feel that, if serious problems exist with the stochastic assumptions, then these problems should be apparent in the residuals. A second conjecture is that less obvious violations in distributional assumptions have negligible effects on estimates and inferences, although Patterson (1999) demonstrates otherwise in an example. Both of these assumptions require verification. We focus on the second assumption in this paper.

The purpose of this paper is to explore the potential sensitivity of SPAs to the assumed distribution of survey indices. A log transformation is commonly used, and is appropriate when the model variability of the survey indices is proportional to the square of their mean (constant coefficient of variation (CV)), although other variance models may be more appropriate for some data sets. We explore SPA estimation for two constant CV stochastic models: the lognormal and the

Received June 28, 2000. Accepted December 21, 2000. Published on the NRC Research Press Web site on March 7, 2001.
J15841

N.G. Cadigan.[1] Fisheries and Oceans, Northwest Atlantic Fisheries Centre, P.O. Box 5667, St. John's, NF A1C 5X1, Canada.
R.A. Myers. Killam Chair of Ocean Studies, Department of Biology, Dalhousie University, Halifax, NS B3H 4J1, Canada.

[1]Corresponding author (e-mail: cadiganN@dfo-mpo.gc.ca).

 DOI: 10.1139/cjfas-58-3-560

gamma. In practice, it is usually difficult to diagnose which of these distributions is the more appropriate one to use (see Atkinson 1982; McCullagh and Nelder 1989). It is often assumed that estimates and inferences are similar for either distribution and that, in practice, it does not matter whether the lognormal or gamma distribution is used. We investigate the bias and precision of maximum likelihood estimates (MLEs) of some SPA parameters for these two distributions using simulation studies. As well, we investigate the bias and precision when the distribution is mis-specified, that is, the change in bias and precision resulting from using lognormal MLEs with gamma-distributed survey indices, and vice versa. We also investigate the accuracy of a method for constructing confidence intervals. Our simulations are based on six real fish populations: Labrador cod, Southern Grand Bank cod, Eastern Scotian shelf cod, Browns Bank cod, Scotian Shelf pollock, and Southern Gulf of St. Lawrence herring.

Materials and methods

SPA can be thought of as “a family of methods for converting catch-at-age data into [population] numbers at age” (Mohn and Cook 1993). Megrey (1989) gives a description of many methods and models used in stock assessments, including SPA. We examine a formulation of SPA that is commonly used for North Atlantic fish stocks (Gavaris 1988; Hilborn and Walters 1992; Mohn and Cook 1993).

We consider a simple cohort model for a population with A age-classes ($a = 1,\ldots,A$) observed over Y years ($y = 1,\ldots,Y$). Let n_{ay} be the unknown number of fish in the population at the beginning of the year. The primary purpose of a SPA is to estimate the survivors in year Y; that is, the $n_{1Y},\ldots,n_{AY}$. The underlying population dynamics model for the SPA formulation we consider is

$$(1)\qquad n_{ay} = n_{a+1y+1}\exp(m) + c_{ay}\exp(m/2)$$

where c_{ay} is the number of age a fish caught by the fishery in year y and m is the annual natural mortality rate. This equation can be used to compute all population numbers (there are $A \times Y$ of these) from the population numbers in the last age and year (there are $A + Y - 1$ of these). We assume that the c and m are known without error. These are standard assumptions and have been used in the assessments of the six stocks we consider in our simulations. Equation 1 is based on the assumption that the commercial catches are taken halfway through the year. This approximation is from Pope (1972) and gives good results; Mertz and Myers (1996) show how eq. 1 can be easily modified to include a seasonal pattern of catches. We denote the survivors up to age A in year Y as the vector $\mathbf{n}_Y$. All n_{Ay}, $y = 1,\ldots,Y-1$, are usually expressed in terms of survivors, using nonlinear constraints (see Gavaris 1988; Myers and Cadigan 1995); however, we simply estimate them. Let $\mathbf{n}_A$ be a vector with elements $n_{A1},\ldots,n_{AY}$. Although not constraining $\mathbf{n}_A$ adds considerably more parameters to be estimated, the estimation is relatively easy, because the population-dynamics model is linear in the parameter vector $\mathbf{n} = [\mathbf{n}_A,\mathbf{n}_Y]$.

Statistical model and estimation

Gamma and lognormal MLEs for the unknown SPA parameters, $\mathbf{n}$, are presented in this section. Survey indices and catches are used to estimate $\mathbf{n}$. We stochastically model the survey indices. In most of the examples we consider, the surveys involve stratified random sampling using a trawl, and the indices are the mean number of fish per tow. It is common to assume that the survey indices reflect only an age-dependent proportion of stock abundance. This is because, among other reasons, the trawl is size-selective. Another common assumption for survey indices is that their SPA model variance is proportional to the square of their mean, that is, constant CV. Using these assumptions, the first two moments of our stochastic SPA model for the random survey indices (R) are

$$(2)\qquad E(R_{ay}) = \mu_{ay} = q_a n_{ay}$$

and

$$(3)\qquad \mathrm{Var}(R_{ay}) = k\mu_{ay}^2$$

where q_a is an age-specific survey-catchability coefficient. Let $\mathbf{q}$ be a vector with elements $q_1,\ldots,q_A$. The CV is $\sqrt{k}$. We explore maximum likelihood estimation of the $\mathbf{n}$ and $\mathbf{q}$ based on two distributions that are commonly used for constant CV stochastic models, namely, the gamma and lognormal distributions (see McCullagh and Nelder 1989).

Gamma SPA MLEs

Consider a gamma-distributed random variable $R > 0$ with $\mu = E(R)$, and let r be a realization of R. The probability density function (PDF) for R is

$$p(r;\mu,\phi) = r^{-1}\Gamma(\phi)^{-1}\left(\frac{\phi r}{\mu}\right)^{\phi}\exp[-\phi r/\mu]$$

where $\phi > 0$. For this distribution $\mathrm{Var}(R) = \mu^2/\phi$, so that $k = 1/\phi$ in eq. 3. The gamma log-likelihood for $\mathbf{n}$, $\mathbf{q}$, and ϕ is

$$\begin{aligned} L_G &= L_G(\mathbf{n},\mathbf{q},\phi\,|\,\{R_{ay} = r_{ay}\}) \\ &= AY\{\phi\log(\phi) - \log\Gamma(\phi)\} \\ &\quad + \phi\sum_{ay}\{\log(r_{ay}/\mu_{ay}) - r_{ay}/\mu_{ay}\} - \sum_{ay}\log r_{ay} \end{aligned}$$

where $\mu_{ay} = \mu_{ay}(\mathbf{n},\mathbf{q})$. MLEs for $\mathbf{n}$ and $\mathbf{q}$ are those that maximize L_G.

It is easy to show that the MLE for q_a when $\mathbf{n}$ is known is

$$\tilde{q}_a = \tilde{q}_a(\mathbf{n}) = Y^{-1}\sum_{y}\frac{r_{ay}}{n_{ay}}$$

Note that $\tilde{q}_a$ is unbiased for q_a. The MLE for ϕ has to be obtained numerically. McCullagh and Nelder (1989) do not recommend using the ϕ MLE for several reasons. They favor using the moment estimator

$$\hat{\phi} = \sum_{ay}\frac{\{(r_{ay} - \mu_{ay})/\mu_{ay}\}^2}{n-p}$$

The estimator of ϕ we use (denoted as $\tilde{\phi}$) is different from $\hat{\phi}$ only in that the $n - p$ term is replaced by n to be consistent with the lognormal variance parameter MLE (see next section). We use $\tilde{q}_a(\mathbf{n})$ and $\tilde{\phi}(\mathbf{n})$, to reduce the number of parameters that must be numerically estimated to just $\mathbf{n}$. We find the MLE for $\mathbf{n}$ using the Newton–Raphson method. The first order derivatives of $L_G(\mathbf{n}, \tilde{\mathbf{q}}, \tilde{\phi}\,|\{R_{ay} = r_{ay}\})$ are not difficult to calculate. We calculate $H(\mathbf{n}) = \partial^2 L_G(\cdot)/\partial\mathbf{n}\,\partial\mathbf{n}'$ by finite differences. Estimates of $\mathrm{Var}(\hat{\mathbf{n}})$ were obtained as the diagonal elements of $-H^{-1}(\hat{\mathbf{n}})$.

Lognormal SPA MLEs

The PDF for the lognormal distribution is

$$p(r;\mu,\sigma^2) = r^{-1}\sigma^{-1}(2\pi)^{-0.5}\exp\{-(\log r - \mu^*)^2/2\sigma^2\}$$

Table 1. Fish populations analyzed.

Population	Unit area	Source	Ages	Years
Labrador cod	2J3KL	Bishop et al. 1993	3–12	1978–1992
Southern Grand Bank cod	3Ps	Bishop et al. 1994	3–12	1978–1993
Eastern Scotian Shelf cod	4VsW	Mohn and MacEachern 1993	3–15	1971–1992
Browns Bank cod	4X	Gavaris et al. 1994	3–10	1970–1992
Scotian Shelf pollock	4VWX	Trippel and Brown 1993	2–12	1974–1992
Southern Gulf of St. Lawrence herring	4T	Claytor et al. 1992	3–10	1978–1991

Note: Ages and years refer to minimum–maximum ages and first–last years used in the sequential population analysis.

Table 2. Number of simulations with large estimates of survivors.

	Survey distribution							
	CV = 0.5				CV = 1			
	Lognormal		Gamma		Lognormal		Gamma	
Stock	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE
3Ps cod	—	—	—	—	—	2	23	2
2J3KL cod	—	—	—	—	—	—	—	—
4VsW cod	—	—	—	—	—	—	—	—
4X cod	1	1	1	1	24	42	61	14
4T herring	—	—	—	—	—	—	27	16
4VWX pollock	—	—	—	—	17	34	42	12
Total	1	1	1	1	41	78	153	44

Note: GMLE and LMLE denote gamma and lognormal maximum likelihood estimators respectively.

where $\sigma > 0$ and $\mu^* = \log(\mu) - \sigma^2/2$. Log($R$) is normally distributed with mean μ^* and variance σ^2. For this distribution,

$$\mathrm{Var}(R) = \exp(\mu^*)^2 \exp(\sigma^2)\{\exp(\sigma^2) - 1\} = k\mu^2$$

where $k = \exp(\sigma^2) - 1$ in eq. 3. The log-likelihood for **n**, **q**, and σ is

$$L_N = L_N(\mathbf{n},\mathbf{q},\sigma^2|\{R_{ay} = r_{ay}\})$$
$$= -AY\log(2\pi\sigma^2)/2$$
$$- \sum_{ay}(\log r_{ay} - \mu^*_{ay})^2 / 2\sigma^2 - \sum_{ay}\log r_{ay}$$

It is easy to show that the MLE for $q_a^* = q_a \exp(-\sigma^2/2)$ when **n** is known is given by

$$\log(\tilde{q}_a^*) = \log\{\tilde{q}_a^*(\mathbf{n})\} = Y^{-1}\sum_y\{\log(r_{ay}) - \log(n_{ay})\}$$

This is an unbiased estimator of $\log(q_a^*)$. Also, the MLE for σ^2 is given by

$$\tilde{\sigma}^2 = \tilde{\sigma}^2(\mathbf{n}) = \frac{\Sigma_{ay}\{\log(r_{ay}) - \log(\tilde{\mu}_{ay})\}^2}{AY}$$

where $\tilde{\mu}_{ay} = \tilde{q}_a^* n_{ay}$. For known **n**, $\tilde{\sigma}^2$ is a biased estimator of σ^2, and the bias is $-\sigma^2/Y$. This bias is of a similar order to the bias in $\hat{\phi}$. It may be more appropriate to use the bias corrected estimator of σ^2; however, this does not affect the estimation of **n**, which is the main focus of this paper. The MLE of q_a can easily be obtained using $\tilde{q}_a^*$ and $\tilde{\sigma}^2$. We use $\tilde{q}_a^*$ and $\tilde{\sigma}^2$ to reduce the number of parameters to estimate, just as for the gamma case. We also use the same method to estimate Var($\hat{\mathbf{n}}$).

Simulations

In this section, we describe the simulation procedures we use to compare gamma and lognormal MLEs of SPA parameters. We consider six fish stocks on which to base our simulations. We use published time series of catch at age and published estimates of n_{aY} ($a = 1,\ldots, A$), n_{Ay} ($y = 1,\ldots,Y - 1$), and q_a, to generate random survey indices at age, using eqs. 1, 2, and 3. We generate 1000 sets of surveys assuming gamma or lognormal errors (2000 in total), and estimate **n** twice for each set using gamma and lognormal MLEs. The only difference between gamma and lognormal simulations is the distribution used to generate the random survey data. We generate three sets of simulated indices, with CVs of 0.25, 0.5, and 1. These values cover the range of SPA CVs that are common in stock assessments. We use the simulations to estimate the bias and other distributional characteristics of the gamma and lognormal MLEs of **n**.

The six fish stocks we consider are all from the Northwest Atlantic (see Table 1). The population parameters are taken from the referenced assessment documents for each stock. We only consider SPAs up to the early 1990s, because commercial fishing moratoria and environmental anomalies after this period complicate the analysis of several of these stocks. We use an annual mortality rate, m, of 0.2, to randomly generate survey data in all our simulation analyses. The same value for m is used in estimation. Documents containing the data used in our analysis can be obtained from the Canadian Stock Assessment Secretariat, Fisheries and Oceans (Station 1256), 200 Kent Street, Ottawa, ON K1A 0E6, Canada, or the Northwest Atlantic Fisheries Organization, P.O. Box 638, Dartmouth, NS B2Y 3Y9, Canada.

Results

In some simulations, particularly when the CV = 1, the parameter estimates were implausibly large. For example, with one of the pollock simulated data sets, the gamma MLE of total survivors converged after 29 iterations and was 3.09×10^{13}; however, after eight iterations, the log-likelihood was within 3.4 units of the maximum but the survivors totalled only 1.6×10^5, which is relatively close to the population total of 9.2×10^4. This is a case where the simulated data

Fig. 1. Standardized biases for lognormal (solid lines) and gamma (dotted lines) maximum likelihood estimators of the sequential population analysis survivors (n_{aY}, $a = 1,\ldots, A$). Panels on the left show biases for lognormal-simulated surveys (coefficient of variation (CV) = 0.5) and panels on the right show biases for gamma-simulated surveys (CV = 0.5). Each row of panels represents a particular fish stock: (*a*) 3Ps cod, (*b*) 2J3KL cod, (*c*) 4VsW cod, (*d*) 4X cod, (*e*) 4T herring, and (*f*) 4VWX pollock.

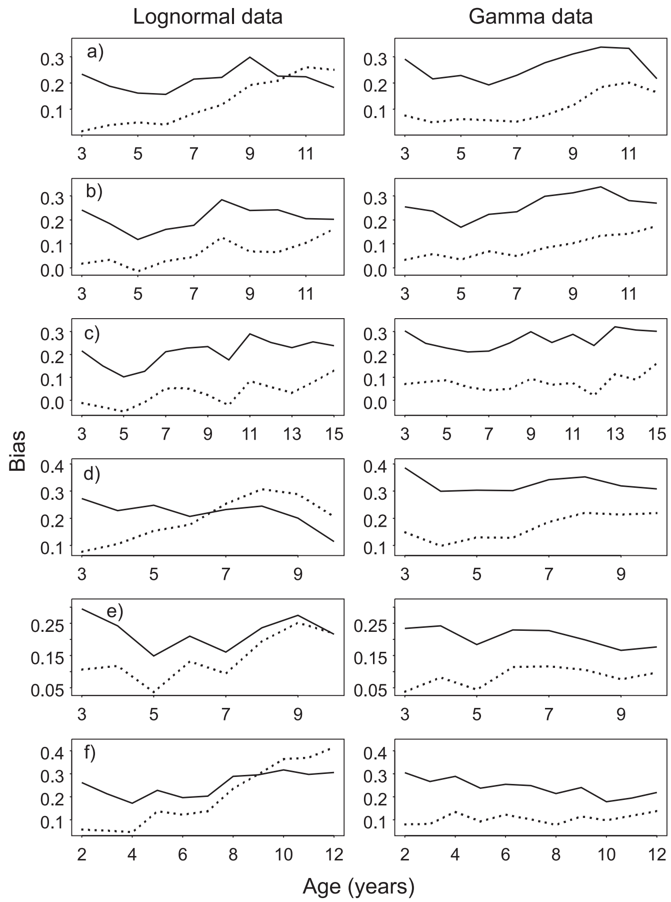

Fig. 2. Standardized biases for lognormal (solid lines) and gamma (dotted lines) maximum likelihood estimators of the sequential population analysis population numbers at the oldest age (n_{Ay}, $y = 1,\ldots,Y$). (*a*) Lognormal-simulated surveys (coefficient of variation (CV) = 0.5). (*b*) Gamma-simulated surveys (CV = 0.5).

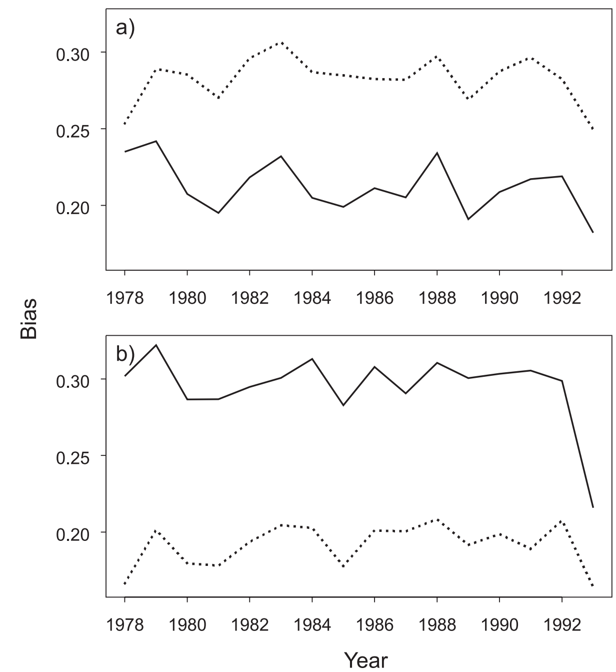

are non-informative about how large the stock might be. We decided to remove these cases, to simulate the process that actually occurs during a stock assessment, that is, such data would not usually be considered in a stock assessment. We removed cases if the estimated total abundance of survivors exceeded 10× the population total abundance of survivors (i.e., $10 \times \Sigma_a n_{aY}$). Both the gamma and lognormal MLEs were removed if either one produced a large estimate of total survivors. The frequency of these large estimates is presented in Table 2. No large estimates were found when the CV = 0.25. When the CV = 1, the trend over all stocks was to obtain more large estimates when the error distribution was wrong, and this occurred more for lognormal MLEs than for gamma MLEs. We examine the consequences of not eliminating the large estimates at the end of this section.

Another problem for the remaining pollock simulations with CV = 1 was that negative variance estimates occurred three times. This happened twice for gamma MLEs and once for lognormal MLEs. Negative variances were always associated with very low estimates of stock size. The corresponding likelihood surfaces were very flat, and these simulations represent data sets that are not informative about how low stock size might be. These cases produced almost singular Hessian matrices that, when inverted, led to incorrect negative variance estimates. These simulations were also removed from further analysis. Note that negative variance estimates were also common in the simulations that produced unreasonably large estimates of stock size.

We focus on the relative performance of gamma and lognormal MLEs of the unknown SPA survivors, $\mathbf{n}_Y$. Let $\hat{\mathbf{n}}_Y^{\mathrm{G}}$ denote the gamma MLEs and let $\hat{\mathbf{n}}_Y^{\mathrm{L}}$ denote the lognormal MLEs. Standardized biases (see eq. 4 below) obtained from the CV = 0.5 simulations for $\hat{\mathbf{n}}_Y^{\mathrm{L}}$ and $\hat{\mathbf{n}}_Y^{\mathrm{G}}$ are presented in Fig. 1. The standardization is the same for both $\hat{\mathbf{n}}_Y^{\mathrm{G}}$ and $\hat{\mathbf{n}}_Y^{\mathrm{L}}$; it is

$$(4) \qquad B_{\mathrm{s}}(\hat{\mathbf{n}}_Y) = \frac{\mathrm{ave}(\hat{\mathbf{n}}_Y) - \mathbf{n}_y}{\mathrm{SD}_{\mathrm{ave}}}$$

where $\mathrm{ave}(\hat{\mathbf{n}}_Y)$ is the average of the simulated $\hat{\mathbf{n}}_Y$s,

$$\mathrm{SD}_{\mathrm{ave}} = [\{\mathrm{Var}(\hat{\mathbf{n}}_Y^{\mathrm{L}}) + \mathrm{Var}(\hat{\mathbf{n}}_Y^{\mathrm{G}})\} / 2]^{1/2},$$

and $\mathrm{Var}(\hat{\mathbf{n}})$ is the variance of $\hat{\mathbf{n}}$ calculated over the simulated estimates. If the only difference in the simulated distributions of $\hat{\mathbf{n}}_Y^{\mathrm{G}}$ and $\hat{\mathbf{n}}_Y^{\mathrm{L}}$ is a location shift, then $\mathrm{SD}_{\mathrm{ave}}$ is simply a combined estimate of the standard deviation. The same $\mathrm{SD}_{\mathrm{ave}}$ is used for both $\hat{\mathbf{n}}_Y^{\mathrm{G}}$ and $\hat{\mathbf{n}}_Y^{\mathrm{L}}$, so that differences in standardized bias are not the consequence of differences in precision; however, a different standardization is used in each panel in Fig. 1 and for each age. We examine the precision of estimators later. Note that the biases in Fig. 1 tend to be midway between the biases we found in the CV = 0.25 and CV = 1 simulations. These results demonstrate that $B_{\mathrm{s}}(\hat{\mathbf{n}}_Y^{\mathrm{G}})$ is often less than $B_{\mathrm{s}}(\hat{\mathbf{n}}_Y^{\mathrm{L}})$, even when the survey distribution is lognormal. This was always the case when the CV = 0.25.

Table 3. Relative bias (%), relative median bias (%), and relative standard deviations (%) for maximum likelihood estimators (MLEs)

	Survey distribution											
	Relative bias (%)								Relative median bias (%)			
	CV = 0.25				CV = 1				CV = 0.25			
	Lognormal		Gamma		Lognormal		Gamma		Lognormal		Gamma	
Stock	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE
3Ps cod (108.6)	2	0	2	1	35	17	98	21	2	–0	2	0
2J3KL cod (335.8)	2	0	3	1	27	6	67	19	1	–1	2	0
4VsW cod (50.6)	1	–0	2	0	21	–1	52	12	1	–1	1	–1
4X cod (21.8)	3	1	3	0	63	88	175	41	1	–1	2	–1
4T herring (1940.6)	3	1	2	0	36	28	136	42	1	0	1	–1

Note: GMLE and LMLE denote gamma and lognormal MLEs. Values in parentheses are population total survivors (millions).

Table 4. Relative bias (%), relative median bias (%), and relative standard deviations (%) for maximum likelihood estimators (MLEs)

	Survey distribution											
	Relative bias (%)								Relative median bias (%)			
	CV = 0.25				CV = 1				CV = 0.25			
	Lognormal		Gamma		Lognormal		Gamma		Lognormal		Gamma	
Stock	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE
3Ps cod (0.10)	–1	1	–1	1	–10	10	–26	4	–2	0	–2	–0
2J3KL cod (0.11)	–1	1	–1	0	–9	13	–23	3	–1	1	–2	–0
4VsW cod (0.46)	–0	1	–1	1	–8	17	–20	5	–1	1	–1	1
4X cod (0.47)	–1	1	–1	2	–15	–1	–30	9	–1	1	–2	1
4T herring (0.07)	–1	0	–0	1	–7	3	–14	7	–1	–0	–1	1

Note: GMLE and LMLE denote gamma and lognormal MLEs. Values in parentheses are the population total exploitation rate.

The biases for n_{Ay} (i.e., $\hat{\mathbf{n}}_A$) when $y < Y$ are similar in magnitude to the bias of $\hat{\mathbf{n}}_{AY}$, which is shown at the right end-point in each of the panels in Fig. 1. We present complete bias results for 3Ps cod in Fig. 2. The results demonstrate that $|B_s(\hat{\mathbf{n}}_A^G)| - |B_s(\hat{\mathbf{n}}_A^L)|$ when the survey distribution is lognormal is less than $|B_s(\hat{\mathbf{n}}_A^L)| - |B_s(\hat{\mathbf{n}}_A^G)|$ when the survey distribution is gamma.

The estimated total abundance of survivors (i.e., $\Sigma_a n_{aY} = n_{+Y}$) is a useful quantity to examine in more detail. In Table 3, we present the bias, median bias, and simulation standard deviations for this quantity. The biases and standard deviations in Table 3 are expressed as percentages of total population (true) abundances (n_{+Y}), which are also shown in this table; for example, the average relative percentage bias is given by

$$B_r(\hat{n}_{+Y}) = 100 \times \left(\frac{\text{ave}(\hat{n}_{+Y}) - n_{+Y}}{n_{+Y}}\right)$$

The median (med) bias is computed by replacing ave(·) with med(·). The results for CV = 0.5 are intermediate between those presented in Table 3.

The relative bias results in Table 3 demonstrate that $|B_r(\hat{n}_{+Y}^G)|$ is usually less than $|B_r(\hat{n}_{+Y}^L)|$. Also in Table 3, we show that the median biases for both $\hat{n}_{+Y}^L$ and $\hat{n}_{+Y}^G$ tend to be less than the average biases, which is to be expected, because the distributions of $\hat{n}_{+Y}^L$ and $\hat{n}_{+Y}^G$ are right-skewed. The median biases demonstrate that, at least 50% of the time, the SPA estimates of n_{+Y} are closer to the population value than the average biases suggest. The median bias of $\hat{n}_{+Y}^G$ tends to be negative quite frequently, whereas the bias of $\hat{n}_{+Y}^L$ is always positive. Hence, we conclude that, more often than not, gamma MLEs will underestimate survivors and lognormal MLEs will overestimate survivors. The relative standard deviation results in Table 3 demonstrate that, when the CV = 0.25, both $\hat{n}_{+Y}^L$ and $\hat{n}_{+Y}^G$ appear to be equally precise, and there is little loss in precision if the nominal MLE distribution is mis-specified. This is not the case when the CV = 1; however, our results show that a greater loss in precision is incurred by using $\hat{\mathbf{n}}_Y^L$ with gamma surveys than by using $\hat{\mathbf{n}}_Y^G$ with lognormal surveys.

Another useful quantity to examine in detail is the total exploitation rate in the last year, that is, $\Sigma_a c_{aY} / \Sigma_a n_{aY} = F_Y$. This is the fraction of the initial stock in year Y removed by the fishery. Exploitation rates are commonly used to assess the sustainability of a fishery. Average and median biases for the total exploitation rate are presented in Table 4. The results when the CV = 1 suggest that, at least 50% of the time, lognormal MLEs underestimate exploitation rates by a considerably greater amount than gamma MLEs do. A disadvantage of gamma MLEs is that the standard deviations of $\hat{F}_Y^G$ (see Table 4) tend to be 40% greater than the standard deviations of $\hat{F}_Y^L$ for lognormal surveys with CV = 1, and 20% greater than the standard deviations of $\hat{F}_Y^L$ for gamma surveys. The effect of this is that for CV = 1, MSE$(\hat{F}_Y^G)$ > MSE$(\hat{F}_Y^L)$ even when the surveys are gamma distributed. This is largely because gamma MLEs overestimate exploitation rates more than lognormal MLEs. For all six stocks and for lognormal data with CV = 1, the average 10th percentiles of $\hat{F}_Y^L - F_Y$ and $\hat{F}_Y^G - F_Y$ are –52 and –49% of F_Y, respectively. For of total survivors.

				Relative standard deviations (%)							
CV= 1				CV = 0.25				CV = 1			
Lognormal		Gamma		Lognormal		Gamma		Lognormal		Gamma	
LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE
21	–1	55	9	12	12	12	12	84	120	216	59
18	–3	48	11	11	11	11	11	52	47	88	54
14	–9	40	5	10	10	10	9	43	40	69	44
26	11	68	9	15	14	15	14	143	356	540	240
15	6	37	6	13	13	13	13	81	83	478	166

of the total exploitation rate in the last sequential population analysis year.

				Relative standard deviations (%)							
CV = 1				CV = 0.25				CV = 1			
Lognormal		Gamma		Lognormal		Gamma		Lognormal		Gamma	
LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE	LMLE	GMLE
–17	1	–35	–9	12	12	11	12	38	54	46	57
–15	4	–32	–10	11	11	11	11	35	51	42	53
–12	9	–28	–4	9	10	9	9	31	45	40	46
–21	–10	–41	–8	14	14	14	14	42	57	49	70
–13	–6	–27	–6	13	13	13	13	43	51	59	64

gamma surveys, the average 10th percentiles of $\hat{F}_Y^{\mathrm{L}} - F_Y$ and $\hat{F}_Y^{\mathrm{G}} - F_Y$ are –68 and –50% of F_Y, respectively. However, the 90th percentiles of $\hat{F}_Y^{\mathrm{L}} - F_Y$ and $\hat{F}_Y^{\mathrm{G}} - F_Y$ are 36 and 70% of F_Y, respectively, for lognormal surveys and 33 and 73% of F_Y, respectively, for gamma surveys.

An important aspect of SPA is the "correctness" of inferences, that is, how well do confidence intervals (CIs) cover population values. Of specific interest to us is how CIs based on the lognormal- or gamma-distribution assumptions compare, especially when the distribution is mis-specified. We investigate this using CIs based on a normal approximation for the distribution of MLEs. Let V^{-1} be the observed information matrix based on ln(**n**) parameters. We nominally assume that

$$(5) \qquad \frac{\ln(\hat{\mathbf{n}}) - \ln(\mathbf{n})}{V^{1/2}}$$

has a standard normal distribution, and use this statistic to construct CIs. These intervals are commonly used for SPA inference, and give a simple way to compare the potential accuracy of inferences for gamma and lognormal MLEs. In our simulations, we count the number of times population values exceed CIs. Lower CIs are of particular importance for fisheries management. In Fig. 3, we plot the observed proportion of times in the CV = 0.5 simulations that population values for survivors were less than the 95% lower confidence limit. The limit based on eq. 5 appears too large, because population values are smaller than this limit with a frequency greater than the nominal 5% value. Nonetheless, the relative inaccuracy of the gamma-based lower limit compared with the lognormal limit when the surveys are lognormal is less than the inaccuracy of the lognormal-based lower limit compared with the gamma limit when the surveys are gamma. Similar results were obtained when the CV = 1. When the CV = 0.25 (results not shown), the limit based on the gamma assumption is more accurate than the limit based on the lognormal assumption, even when the distribution is mis-specified.

In the above analyses, we eliminated estimates that were implausibly high, that is, 10 times the truth. Our general results were robust to changes to this procedure. If fewer very large estimates were removed, the bias of the lognormal estimator increased greatly, while if estimates five times the truth were eliminated, the bias was reduced. In both cases, the gamma MLEs seemed preferable.

Discussion

We have found that assumptions about the distribution of research-survey indices in the method most commonly used to estimate the abundance of marine fish populations in the North Atlantic have a large effect on the bias of important parameters for fisheries managers. The lognormal distribution is commonly used for SPA estimation and inference; however, the gamma distribution appears preferable. This is because gamma MLEs have less bias even when the true distribution of survey indices is lognormal or, at least, the increase in bias when using gamma MLEs with lognormal-distributed surveys is less than the increase in bias when using lognormal MLEs with gamma-distributed surveys. Our results also suggest that, for estimating population size, the loss in precision using

Fig. 3. Simulation probabilities that sequential population analysis survivors (n_{aY}, $a = 1,\ldots, A$) are less than lower confidence limits based on lognormal (solid lines) and gamma (dotted lines) maximum likelihood estimators. Panels on the left show biases for lognormal-simulated surveys (coefficient of variation (CV) = 0.5). Panels on the right show biases for gamma-simulated surveys (CV = 0.5). Each row of panels represents a particular fish stock: (*a*) 3Ps cod, (*b*) 2J3KL cod, (*c*) 4VsW cod, (*d*) 4X cod, (*e*) 4T herring, and (*f*) 4VWX pollock. The nominal probability of exceeding the confidence limit is shown as a broken line.

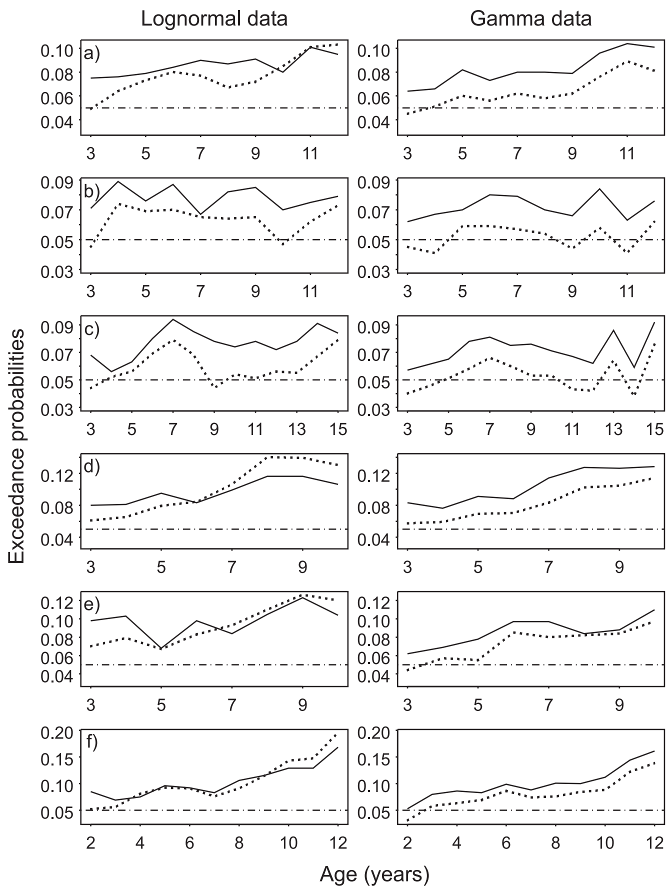

lognormal MLEs with gamma-distributed surveys tends to be greater than the loss in precision using gamma MLEs with lognormal-distributed surveys. Our results also suggest that more accurate and robust lower CIs for population size may be obtained using the gamma-distribution assumption rather than the lognormal-distribution assumption.

Our conclusions agree with those presented by Firth (1988), who showed that the loss in asymptotic efficiency of gamma MLEs when errors are lognormally distributed is less than the loss associated with lognormal MLEs when the errors are gamma distributed. Our results are also consistent with Myers and Pepin (1990), who found that methods that rely on the lognormal assumption may produce large biases and loss of efficiency, if there are even small deviations from the assumptions.

One deficiency of gamma MLEs of SPA exploitation rates compared with lognormal MLEs is that gamma MLEs were found to be less precise than lognormal MLEs, even when the surveys were gamma distributed. This could be serious, because exploitation rates often play an important role in setting commercial-fishery catch quotas; that is, often quotas are selected that give a low predicted exploitation rate. Our results show that the difference in precision is because gamma MLEs tend to overestimate exploitation rates more than lognormal MLEs, that is, the probability of getting a large estimate of the exploitation rate in the last SPA year is considerably greater for gamma MLEs than for lognormal MLEs. This is a less serious conservation error than underestimating exploitation rates and overestimating stock size, which is more of a problem for lognormal than gamma MLEs. Consequently, the decreased precision of gamma MLEs for estimating exploitation rates compared with lognormal MLEs is not a good reason to favor the use of lognormal MLEs.

Some of our assumptions may be violated for some data sets. In particular, there may be errors in the observed commercial catches, and natural mortality may vary over time and age. Violation of either of these assumptions may create biases and increase estimation error variance (Lapointe et al. 1989; Clark 1999). Trends in unreported catch or discarding can have disastrous consequences in a collapsing fishery (Myers et al. 1997). Another problem is that survey indices may not be independent within years (Myers and Cadigan 1995) for some stocks. We have limited our analyses to the model defined by eqs. 2 and 3, because it accounts for a major source of variability (surveys, see Rivard 1989), and it is desirable for our analyses to be applicable to fish stock assessment methods that are commonly used, at least for North Atlantic stocks. It also seems unlikely that lognormal MLEs may be preferable to gamma MLEs, because some other part of the SPA model is mis-specified. Nonetheless, the estimation biases we report may be small for some stocks compared with the biases that can result from using incorrect catches and wrong values for natural mortality.

Our approach to SPA estimation is somewhat nonstandard in North Atlantic assessments, in that we treat as unknown both the survivors ($\mathbf{n}_Y$) and the numbers in the oldest age-class of the cohort model ($\mathbf{n}_A$). Typically, the latter numbers are approximated using constraints on their fishing mortalities. This reduces the number of parameters to estimate, but also complicates estimation. We feel that our results will apply when constraints on fishing mortalities are used, providing the constraints are good approximations of reality. Otherwise the choice of error distribution may be a minor problem in terms of the biases caused by incorrectly modeling the $\mathbf{n}_A$. In this respect, our results also suggest that, when the survey SPA errors are not too large (CV ≤ 0.5), then it is possible to estimate the $\mathbf{n}_A$ and $\mathbf{n}_Y$ with reasonable accuracy, and we recommend doing this at least as a diagnostic check for the suitability of fishing-mortality constraints.

More complicated and possibly more realistic stochastic assumptions are also used for modeling catch-at-age data. Examples are found in Deriso et al. (1985) and Gudmundsson (1994). These methods all involve stochastic models for commercial catches. The catch models usually express c_{ay} in terms of population numbers, using a commercial fishery selectivity model. These methods may also allow for errors in assumptions about the annual natural mortality rate, *m*. Quite often these methods incorporate additional information, like the effort expended by the commercial fleet to obtain catches. The applicability of our results to these other methods is difficult to predict, however our comments in the previous two paragraphs also apply to these more complicated approaches; for example, if the fishery-selectivity model is mis-specified, then the choice of survey SPA error distribution may be a minor problem compared with the effect of incorrectly modeling commercial-catch information.

Gamma and lognormal MLEs produce different estimates, essentially because both distributions give different weights to the deviations between observations and predictions. An alternative approach is to use quasi-likelihood methods (see McCullagh and Nelder 1989), as they only use information on the first two moments.

Acknowledgement

The authors thank C. Bishop, B. Brodie, G. Chouinard, S. Gavaris, G.A. Nielsen, and D. Stansbury of the Canadian Department of Fisheries and Oceans for their data.

References

Atkinson, A.C. 1982. Regression diagnostics, transformations, and constructed variables. J. Roy. Statist. Soc. Ser. B, **44**: 1–36.

Bishop, C.A., Murphy, E.F., Davis, M.B., Baird, J.W., and Rose, G.A. 1993. An assessment of the cod stock in NAFO Div. 2J3KL. NAFO (Northwest Atlantic Fisheries Organization) SCR Doc. 93/86. Ser. No. N2271. NAFO, P.O. Box 638, Dartmouth, NS B2Y 3X9, Canada.

Bishop, C.A., Murphy, E.F., and Davis, M.B. 1994. An assessment of the cod stock in NAFO subdivision 3Ps. DFO (Department of Fisheries and Oceans) Atl. Fish. Res. Doc. 94/33. Canadian Stock Assessment Secretariat, Fisheries and Oceans Canada (Station 1256), 200 Kent St., Ottawa, ON K1A OE6, Canada.

Clark, W.G. 1999. Effects of an erroneous natural mortality rate on a simple age-structured stock assessment. Can. J. Fish. Aquat. Sci. **56**: 1721–1731.

Claytor, R.R., Nielsen, G., Dupuis, H.M.C., and Mowbray, F. 1992. Assessment of Atlantic herring in NAFO division 4T, 1991. CAFSAC (Canadian Atlantic Fisheries Scientific Advisory Committee) Res. Doc. 92/76. Canadian Stock Assessment Secretariat, Fisheries and Oceans Canada (Station 1256), 200 Kent St., Ottawa, ON K1A OE6, Canada.

Cook, R.M., Kunzlik, P.A., and Fryer, R.J. 1991. On the quality of North Sea cod stock forecasts. ICES (Int. Counc. Explor. Sea) J. Mar. Sci. **48**: 1–13.

Cooke, J.G. 1985. On the estimation of trends in year class strength using cohort models. Rep. Int. Whal. Comm. **35**: 325–330.

Deriso, R.B., Quinn, T.J., and Neal, P.R. 1985. Catch-age analysis with auxiliary information. Can. J. Fish. Aquat. Sci. **42**: 815–824.

Firth, D. 1988. Multiplicative errors: log-normal or gamma? J. Roy. Statist. Soc. Ser. B, **50**: 266–268.

Gavaris, S. 1988. An adaptive framework for the estimation of population size. CAFSAC (Canadian Atlantic Fisheries Scientific Advisory Committee) Res. Doc. 88/29. Canadian Stock Assessment Secretariat, Fisheries and Oceans Canada (Station 1256), 200 Kent St., Ottawa, ON K1A OE6, Canada.

Gavaris, S., Clark, D., and Hammel, J. 1994. Assessment of cod in Division 4X. DFO (Department of Fisheries and Oceans) Atl. Fish. Res. Doc. 94/36. Canadian Stock Assessment Secretariat, Fisheries and Oceans Canada (Station 1256), 200 Kent St., Ottawa, ON K1A OE6, Canada.

Gudmundsson, G. 1994. Time series analysis of catch-at-age observations. Appl. Statist. **43**: 117–126.

Hilborn, R., and Walters, C.J. 1992. Quantitative fisheries stock assessment: choice, dynamics and uncertainty. Chapman and Hall, New York.

ICES (International Council for the Exploration of the Sea). 1991. Report of the working group on methods of fish stock assessments. ICES (Int. Counc. Explor. Sea) CM 1991/Assess: 25. ICES, Palaegade 2-4, DK-1261, Copenhagen K, Denmark.

Lapointe, M.F., Peterman, R.M., and MacCall, A.D. 1989. Trends in fishing mortality rate along with errors in natural mortality rate can cause spurious time trends in fish stock abundances estimated by virtual population analysis (VPA). Can. J. Fish. Aquat. Sci. **46**: 2129–2139.

McCullagh, P., and Nelder, J.A. 1989. Generalized linear models. Chapman and Hall, New York.

McCullough, D.R. 1979. The Georges Reserve deer herd. The University of Michigan Press, Ann Arbor.

Megrey, B.A. 1989. Review and comparison of age-structured stock assessment models from theoretical and applied points of view. Am. Fish. Soc. Symp. **6**: 8–48.

Mertz, G., and Myers, R.A. 1996. An extended cohort analysis: incorporating the effect of seasonal catches. Can. J. Fish. Aquat. Sci. **53**: 159–163.

Mohn, R.K., and Cook, R. 1993. Introduction to sequential population analysis. NAFO (Northwest Atl. Fish. Organ.) Sci. Counc. Stud. No. 17.

Mohn, R.K., and MacEachern, W.J. 1993. Assessment of 4VsW cod in 1992. DFO (Department of Fisheries and Oceans) Atl. Fish. Res. Doc. 93/22. Canadian Stock Assessment Secretariat, Fisheries and Oceans Canada (Station 1256), 200 Kent St., Ottawa, ON K1A OE6, Canada.

Myers, R.A., and Cadigan, N.G. 1995. Statistical analysis of catch-at-age data with correlated errors. Can. J. Fish. Aquat. Sci. **52**: 1265–1273.

Myers, R.A., and Pepin, P. 1990. The robustness of lognormal-based estimators of abundance. Biometrics, **46**: 1185–1192.

Myers, R.A., Hutchings, J.A., and Barrowman, N.J. 1997. Why do fish stocks collapse? The example of cod in Eastern Canada. Ecol. Appl. **7**: 91–106.

Patterson, K.R. 1999. Evaluating uncertainty in harvest control law catches using Bayesian Markov chain Monte Carlo virtual population analysis with adaptive rejection sampling and including structural uncertainty. Can. J. Fish. Aquat. Sci. **56**: 208–221.

Pope, J.G. 1972. An investigation of the accuracy of virtual population analysis using cohort analysis. Int. Comm. Northwest Atl. Fish. Res. Bull. **56**: 65–74.

Rivard, D. 1989. Overview of the systematic, structural, and sampling errors in cohort analysis. Am. Fish. Soc. Symp. **6**: 49–65.

Roff, D.A., and Bowen, W.D. 1983. Population dynamics and management of the Northwest Atlantic harp seal (*Phoca groenlandica*). Can. J. Fish. Aquat. Sci. **40**: 919–932.

Shepherd, J.G. 1988. Fish stock assessments and their data requirements. *In* Fish population dynamics. *Edited by* J.A. Gulland. Wiley, New York. pp. 35–62.

Trippel, E.A., and Brown, L.L. 1993. Assessment of Pollock (*Pollachius virens*) in Divisions 4VWX and Subdivision 5Zc for 1992. DFO (Department of Fisheries and Oceans) Atl. Fish. Res. Doc. 93/60. Canadian Stock Assessment Secretariat, Fisheries and Oceans Canada (Station 1256), 200 Kent St., Ottawa, ON K1A OE6, Canada.