



Survey estimates of king crab (*Paralithodes camtschaticus*) abundance off Northern Norway using GLMs within a mixed generalized gamma-binomial model and Bayesian inference

C. Hvingel*, M. C. S. Kingsley, and J. H. Sundet

Havforskningsinstituttet – Institute of Marine Research, PO Box 6404, N-9294 Tromsø, Norway

*Corresponding Author: tel: +47 77 60 97 50; fax: +47 77 60 97 01; e-mail: carsten.hvingel@imr.no

Hvingel, C., Kingsley, M.C.S., and Sundet, J.H. 2012. Survey estimates of king crab (*Paralithodes camtschaticus*) abundance off Northern Norway using GLMs within a mixed generalized gamma-binomial model and Bayesian inference. – ICES Journal of Marine Science, 69: 1416–1426.

Received 1 July 2011; accepted 30 May 2012; advance access publication 5 July 2012.

A trawl survey provides information on number and biomass of introduced king crab (*Paralithodes camtschaticus*) to the management of a fishery off the coast of Northern Norway; the annual catch quotas are largely set as a percentage of the survey estimate. A specially built sledge trawl was designed for the survey. It needs only small areas of trawlable bottom, performs well on a wide range of bottoms, and appears to have good catchability for benthic organisms. Many survey hauls catch no crabs and the non-zero catches have a highly skewed distribution. Data were therefore analysed with a compound model, in which separate predictors were fitted for the proportion of zero catches and for the catch size of the non-zero catches. The compound model was fitted by Bayesian methods using WinBUGS. The distribution of non-zero catches fitted well to a generalized gamma distribution, but with parameter values that made it approximate a lognormal distribution. Numbers of fishable crabs peaked in 2003, and total numbers in 2010 were about two-fifths of the 2003 maximum.

Keywords: Bayesian, delta model, generalized gamma, king crab (*Paralithodes camtschaticus*), survey, zero catch.

Introduction

The king crab (*Paralithodes camtschaticus*) is an introduced species in the Barents Sea. Crabs were released into the eastern Barents Sea in the 1960s and the species' abundance has since increased to where it can support a coastal fishery off Northern Norway; close to 2000 tons of crabs were taken in 2008. The size and condition of the stock are monitored by survey, and within a restricted area the fishery is managed by catch control; annual allowable catches are set as a percentage of the overall survey biomass estimates. Similar types of harvest-rate based harvest-control rules have been used for stocks of this species in the Bering Sea (Otto, 1986; Zheng *et al.*, 1997). Other biological information, e.g. stock demographics (recruitment, sex- and length-distribution), migration patterns, and interaction with other species, is also considered in the stock assessments, but estimates of numbers and biomass obtained from the annual surveys are central to the current management framework.

The trawl surveys face two major challenges: (i) rough bottoms, which hinder trawling operations in many areas inhabited by

crabs, and (ii) analysing the data, which include a large proportion of empty hauls. The methods chosen to accommodate these challenges may affect the precision and bias of the estimates of stock abundance, and thus the management of the fishery.

Owing to the rough bottoms often found within the range of the Red King Crab in the Norwegian Barents Sea, standard (otter) trawling gears, as used for crab surveys in Alaska (Otto, 1986) and Russia (Berenboim and Pinchukov, 2005), were not considered suitable. Using such gear would significantly reduce survey coverage, impose unknown biases on stock-size estimates, and make them sensitive to changes in distribution if crabs moved between trawlable and non-trawlable areas. A special “crab trawl” was constructed and is now used in the survey. We describe the design of the trawling gear and how the Norwegian king crab survey is conducted.

A prevalent feature of the Red King Crab survey data is “hauls without catch”. These “zeros” appear in most fish stock survey data—if not always at the level of total catch then at least when the data are treated at finer demographic resolutions (age, length or stage groups). It is a common problem (e.g. Maunder and

Punt, 2004; Fletcher *et al.*, 2005; Martin *et al.*, 2005), and data such as this—highly skewed, and not drawn from a standard distribution—pose problems in analysis, in particular with respect to the calculation of confidence distributions, even for the estimate of a simple statistic such as total numbers.

The following methods for analysing data containing zeros have been used, or recommended, from time to time: (i) Aggregate data over areas and/or stock demography until zeros disappear, make the estimates and then segregate the result according to the measured average demographic structure. This may entail assuming, e.g., that the spatial distribution of different life-stages is similar, which is not necessarily the case. (ii) Log-transform the data: add a positive constant to get rid of the zeros; log-transform these “new” data, average and sum, then back-transform the summary statistics by exponentiating and subtracting the constant—as described in popular statistical textbooks (e.g. Sokal and Rohlf, 1995, p. 421). This method has been supported by the observation that, if data are really drawn from a lognormal distribution, log-transforming, analysing and then back-transforming can improve estimates of population parameters. However, the authors and others (e.g. McArdle and Anderson, 2004; Fletcher *et al.*, 2005; Martin *et al.*, 2005) do not recommend this method. The back-transformation of the summary statistics usually involves approximations, if not downright errors, and the results obtained depend on the constant chosen and the scaling of the data (e.g. scaling hauls to m^2 , km^2 or other before the calculations are performed); this procedure can lead to large errors. (iii) Simply ignore the zeros: this will lead to bias and to a large underestimation of the uncertainty surrounding the result.

Other methods providing a statistically more rigorous treatment of survey data including zeros have been developed (e.g. Maunder and Punt, 2004; Fletcher *et al.*, 2005; Martin *et al.*, 2005; Shono *et al.*, 2008; Zuur *et al.*, 2009). The basic “delta”

method comprises the fitting of two separate models, one for occurrence (zeros and non-zeros) and a second for the distribution of densities at sites with non-zero catches, either a lognormal or a gamma distribution (Pennington, 1983, 1996; Myers and Pepin, 1990). Stefánsson (1996) suggests using a multi-year GLM within the delta approach, and this combination has the advantage that by analysing an entire series, instead of just one year at a time, the spatial pattern of stock density can be modelled by including area variables. Further structuring—e.g. environmental—variables can also be added to account for components of the residual variation other than a year effect. Such a GLM structure can also accommodate missing data, typically originating from incomplete surveys.

We present here a Bayesian version of a GLM within the delta model, and generate probability density distributions of the annual survey estimates in the units of the raw data. It has been possible to analyse the non-zero data without assuming that it necessarily has either a lognormal or a standard gamma distribution.

Method

Survey gear and design

Since 2000 standardized bottom-trawl surveys have provided annual indices of abundance, distribution and biological information for the stock of king crab inhabiting waters off Northern Norway. The survey area comprises four large fjords in the north coast (Figure 1), each considered as a stratum. The area strata are further divided in three depth strata: 0–90 m, 90–330 m and >330 m. Most of the survey area is less than 330 m deep, but about 7% of the Varanger fjord stratum, the easternmost, is deeper than 330 m (Table 1). The survey uses a fixed-station design in which the same locations are trawled each year. Stations were selected to provide coverage of all areas and

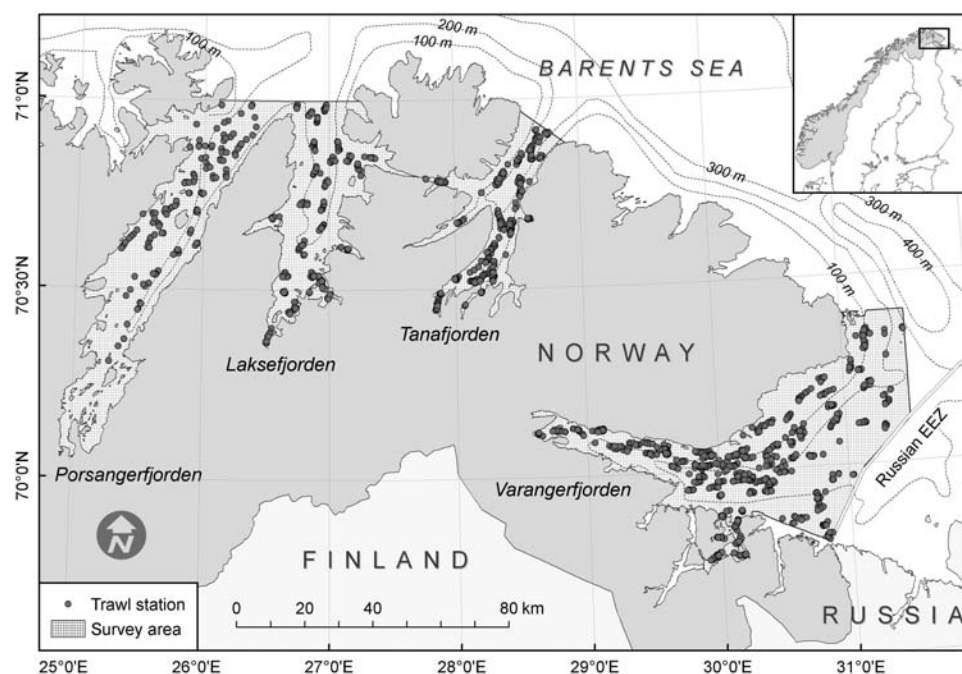


Figure 1. King Crab in northern Norway: survey area consisting of the four fjords Varanger, Tana, Lakse and Porsanger. Dots indicate the location of hauls.

Table 1. King crabs in Northern Norway: trawl survey area (km²) by fjord and depth class

Depth	Varangerfjord	Tanaufjord	Laksefjord	Porsangerfjord	Total
< 90 m	617	228	418	473	1 736
90–330 m	2 053	535	801	859	4 248
> 330 m	198	0	0	0	198
Total	2 868	763	1 219	1 332	6 182

depths. From 2000 to 2006 the area surveyed expanded in keeping with the development and distribution of the crab stock (Hjelset et al., 2009). The number of hauls increased from 87 in 2000 to 165 in 2006, but has averaged only 123 in 2008–2010.

The gear used is a specially constructed version of a sledge (Agassiz-type) trawl without doors or wings. The mouth of the “crab trawl” is shaped by a rectangular steel frame 6 m wide and 1 m high; a 15 m long net is mounted behind. Each side of the heavy frame is supported by 20 cm wide steel runners to prevent it from sinking into soft mud. The frame is buoyed by 50 trawl buoys, each with a buoyancy of 8.5 kg. The lower part (bottom panel) of the trawl is made of double Polar Goll® 5 mm twine with a mesh size of 135 mm, while the upper part (roof panel) consists of PE-net® with a mesh size of 40 mm. In the cod-end PE-net® is used as an inner net covered by Polar Goll®. Four trawl buoys with a buoyancy of 4 kg each are mounted along each side of the cod-end to lift it off the bottom and reduce damage. The trawl is towed by four wires, two attached to the frame at each side.


The survey is conducted in August and September when the crabs are not aggregated for mating but instead are more evenly distributed. The trawl is towed at 1.5 knots for 30 min from the time the trawl contacts the bottom until the start of haul-back. The area swept by a standard haul is reckoned to be 8334 m². Crabs caught in the trawl are measured and sexed. The statistic of greatest interest is the number of crabs over 137 mm CL (carapax length), as it is this size class that supports the fishery, and for that reason the analysis presented here is concerned with this size class.

Distribution of the data

The data are considerably skewed (skewness 8.65) and include more zeros than can be fitted by standard distributions (Figure 2). The many zeros reflect the typical patchy distribution of this species, but also the extent of its spatial distribution within the surveyed areas. Overall stock size is a conflation of the extent of its distribution (the fraction of non-zeros) and the average density at occupied sites (the non-zeros); we modelled these two processes separately.

The probabilities of zero catch were modelled by a binomial distribution. The non-zero part of the data is also positively skewed (skewness = 7.04). The gamma or lognormal distributions are typically used to model such data (Myers and Pepin, 1990; Stefánson, 1996; Pennington, 1996), and discussions on which of the two to use have been frequent (Syrjala, 2000; Dick, 2004). However, several other distributions might also be candidates. We tested different distributions using goodness of fit statistics (Kolmogorov-Smirnov, Anderson-Darling, chi-squared) and visual inspections of probability plots across the data matrix to provide some guidance as to which to prefer.

Different approaches were now open, e.g.: (i) Run all models (each with its different assumptions about the distribution of

**Figure 2.** King Crab in northern Norway: distribution of number of crabs larger than 137mm CL in individual survey hauls.

the data) one by one and pick the “true” one judged by a relevant score of best fit statistics. (ii) Run all models simultaneously and average the results in order to include all the different hypotheses of data distribution. Option (i) seemed a bit restrictive and would likely trap us in the old lognormal vs. gamma discussion. While option (ii) was closer to ideal by including model uncertainty, the problem of how to weigh the different models in the averaging process still needed to be resolved, and the accompanied increase in model running time due to the increased complexity made this approach impractical. As a pragmatic solution we decided to use a generalized gamma distribution as the model for the positive data. It has the great advantage of being a three-parameter distribution, and therefore offers at least the possibility of being able to fit mean, variance and skewness with some measure of independence. The generalized gamma includes a number of standard non-negative, positively skewed distributions as special cases (e.g. the lognormal, gamma, exponential, Weibull etc.). Loglogistic and inverse Gaussian distributions are not included among them and were therefore not investigated further, although the initial analyses had identified them as possible candidates. We also fitted a lognormal model and, further, compared these results with those obtained by simply taking an arithmetic mean of all stations in each sampling unit (a “sampling unit” is one depth class, in one fjord, in one year).

The generalized gamma distribution

The generalized gamma distribution provides a general approach to the problem of fitting a distribution to skewed data. Its probability density function in the form commonly used (Stacy, 1962; Stacy and Mihram, 1965) is:

$$f(x) = \frac{\beta}{\Gamma(\kappa)\theta} \left(\frac{x}{\theta}\right)^{\kappa\beta-1} \exp\left(-\left(\frac{x}{\theta}\right)^\beta\right)$$

where $\theta (>0)$ is a scale parameter, $\beta (>0)$ and $\kappa (>0)$ are shape parameters and $\Gamma(\kappa)$ is the gamma function of κ . Moments of the distribution are given by:

$$E(x') = \theta \frac{\Gamma\left(\kappa + \frac{r}{\beta}\right)}{\Gamma(\kappa)}.$$

leading to

$$\theta \frac{\Gamma\left(\kappa + \frac{1}{\beta}\right)}{\Gamma(\kappa)}$$

for the mean of the distribution, and

$$\sqrt{\frac{\Gamma\left(\kappa + \frac{2}{\beta}\right)\Gamma(\kappa)}{\left(\Gamma\left(\kappa + \frac{1}{\beta}\right)\right)^2} - 1}$$

for its CV.

In this form the distribution has a long-standing reputation for being difficult to fit: the parameters do not represent recognizable properties of the distribution, so it is difficult to provide starting values for fitting methods, and they are strongly related so that convergence can be difficult to achieve.

Prentice (1974), working with the logarithm of the argument of the generalized gamma, mapped the lognormal limiting distribution, as k tends to ∞ , to the origin ($l=0$) using the transformation $\lambda = 1/\sqrt{k}$ and gave a limiting expression for the variance of the transformed variable (approximately the CV^2 of the parent) as $\sigma^2 = 1/\beta^2\kappa$.

The logarithm of the mean is approximately

$$\mu = \ln(\theta) + \ln\left(\Gamma\left(\kappa + \frac{1}{\beta}\right)\right) - \ln(\Gamma(\kappa))$$

which for β not small can be approximated by

$$\mu = \ln(\theta) + \frac{1}{\beta} \ln(k)$$

Prentice (loc. cit.) also demonstrated that the logged distribution, with transformed parameters and parameter space extended to include negative l (so that the centre of the parameter space was occupied by a symmetrical normal distribution), was somewhat easier to fit by maximum likelihood (ML) methods than the standard form of the generalized gamma distribution.

A main effects model for a survey-based stock size index

The distribution of the number of individuals (y) in a haul may be approximated by a mixed probability function consisting of a binary discrete distribution for zero/non-zero and a continuous distribution, $g(y)$, for the non-zero values:

$$f(y) = \begin{cases} 1 - \pi & y = 0 \\ \pi \cdot g(y) & y > 0 \end{cases}$$

where π is the probability of a positive haul and $\int_0^{\infty} g(y)dy = 1$. The expectation of y , i.e. $E(y)$, equal to $\pi\bar{g}$, i.e. the probability of getting a positive haul times the mean catch in the positive hauls, is the survey index of mean number for which an estimate is sought;

using Bayesian methods for fitting the model made confidence intervals for these indices more accessible. Two separate General Linear Models (GLMs) can then be constructed to predict, respectively, π and $g(y)$ and define the overall distribution of the data. In this example we use area (the four fjords, Figure 1) and depth (three ranges: $<90\text{m}$, $90-330\text{m}$, $>330\text{m}$) as explanatory variables taking spatial structure into account, and a time effect (year) as surveys are conducted on an annual basis.

The GLM model for the occurrence of non-zero hauls

Of n_{dat} hauls taken in depth range d in area a in the year t , m_{dat} catch crabs. The variable m_{dat} is considered to have a binomial distribution with probability parameter π_{dat} :

$$m_{dat} \sim \text{Binomial}(n_{dat}, \pi_{dat}).$$

Using a logit link function to scale probabilities to effects, the probability π was described as having components:

$$\text{logit}(\pi_{dat}) = \alpha_{1,ta} + \alpha_{2,d}$$

where $\alpha_{1,at}$ was an effect for area a in year t and $\alpha_{2,d}$ was the effect of depth stratum d . The invasion of Northern Norway waters by crabs was progressive, from the east, so each year-area combination had to be modelled as a separate effect; but within each area the different depths were considered to be invaded synchronously, so depth was modelled as completely crossed with year-area combinations.

The GLM model for the abundance of crabs in non-zero hauls

The density, y_{itad} , in the haul i taken at depth d in area a in the year t with positive catch was modelled by a generalized gamma distribution. This model was kept simple by assuming that the distribution of densities had the same shape in all sampling units, i.e. that the CV of density and its skewness were the same, regardless of the mean. In marine resource surveys, it is common to find that the CV of density is the same in sampling units of very different mean density. Furthermore, of the 68 sampling units with any non-zero catches, 25% only had one catch, so no information on CV, and 40% one or two, so no information on skewness. It was impracticable to fit CV and skewness separately to each sampling unit. This permitted fitting common values of kappa and beta, but different values for the mean density, to the year-area-depth combinations, so:

$$y_{itad} \sim \text{gen.gamma}(\kappa, \theta_{itad}, \beta).$$

Using a log link the predicted logarithm of mean density is:

$$\mu_{itad} = \gamma_{1,ta} + \gamma_{2,d}$$

where $\gamma_{1,ta}$ is an effect of area a in year t , and $\gamma_{2,d}$ is the effect of depth stratum d ; as for the model for occurrence of non-zero hauls, and for the same reason, year-area combinations were considered separately but crossed with depth effects.

In applying Bayes' equation to the present problem, the posterior probability distribution of the survey indices was derived by Monte-Carlo-Markov-Chain (MCMC) sampling methods (see e.g. Congdon, 2001). The programming framework WinBUGS 1.4 provided a means of specifying and analysing a Bayesian

model, including selection and implementation of appropriate algorithms (Spiegelhalter *et al.*, 2004).

Lacking prior knowledge on the density or prevalence of the species, we used uninformative prior distributions for the model parameters. The year-area and depth effects, for both presence and density, were given uniform priors (in log space) between -10 and 10 , a range much wider than that of the observed densities. The (approximate) CV of the generalized gamma distribution, σ , was given a uniform prior between 0.5 and 4 . The shape parameter, λ , was given a distribution uniform in log space between -4 and the logarithm of the reciprocal of σ . This upper limit was chosen to ensure that the distribution of non-zero densities was not modelled as J-shaped, while the lower limit avoided computational difficulties that occurred in sampling from the generalized gamma distribution with large values of κ (which is the reciprocal of λ^2).

While the generalized-gamma model was flexibly able to fit distributions ranging from lognormal to gamma under an

assumption of uniform CV in all sampling units, its complex parameterization made it cumbersome to fit with variable CVs. Examination of the data showed that in these surveys, the CV of density in non-empty hauls was close to proportional to the sampling-unit mean (Figure 3), but we were unable to parameterize the generalized gamma model to fit such data. Therefore, we also fitted a lognormal distribution defined by $\log(y_{itad}) \sim N(m_{itad}, s_{itad}^2)$ where the mean of the lognormal, $\exp(m_{itad} + s_{itad}^2/2) = \exp(\mu_{itad})$ and $\mu_{itad} = \gamma_{1,ta} + \gamma_{2,d}$. The relationship between CV and mean density was then modelled by settings_{itad} = $\eta_d \cdot \exp(\mu_{itad})$. TA year-area means γ_1 , D-1 depth effects γ_2 , and D coefficients η were then fitted for the relationship between CV and mean. The year-area means and the depth effects, parameters in log space, were given prior distributions uniform from -10 to 10 (as when fitting the generalized gamma), and the coefficients η were given priors uniform in log space between -3 and 3 .

The survey indices

For each year-area-depth combination the proportion of the area occupied by crabs was estimated by reversing the logit scaling:

$$P_{ta,d} = \frac{\exp(\alpha_{1,ta} + \alpha_{2,d})}{1 + \exp(\alpha_{1,ta} + \alpha_{2,d})}$$

and the mean density of crabs in the occupied area was estimated by

$$M_{ta,d} = \exp(\gamma_{1,ta} + \gamma_{2,d})$$

so the mean density of crabs in the area-depth combination ad in year t was $P_{ta,d} \cdot M_{ta,d}$, the number was $A_{a,d} \cdot P_{ta,d} \cdot M_{ta,d}$ and the total number in area a was $\sum_d A_{a,d} P_{ta,d} M_{ta,d}$

where $A_{a,d}$ was the extent of area ad .

Results

Model fit and diagnostics

The MCMC process of the generalized gamma-binomial model converged orderly, and the model fitted the data well. All priors are strongly updated (Figure 4) and generally not constrained by limits on prior distributions; the exception was the limit of -4 on the distribution of $\log(\lambda)$ in fitting the generalized gamma,




Figure 3. King Crab in northern Norway: sampling-unit CV of density in non-empty hauls vs mean density.




Figure 4. King Crabs in northern Norway: prior and posterior cumulative probability density distributions of the parameters σ and λ of the generalised gamma model (A and B) and for the parameter a which relates the mean and CV of the lognormal model (C).

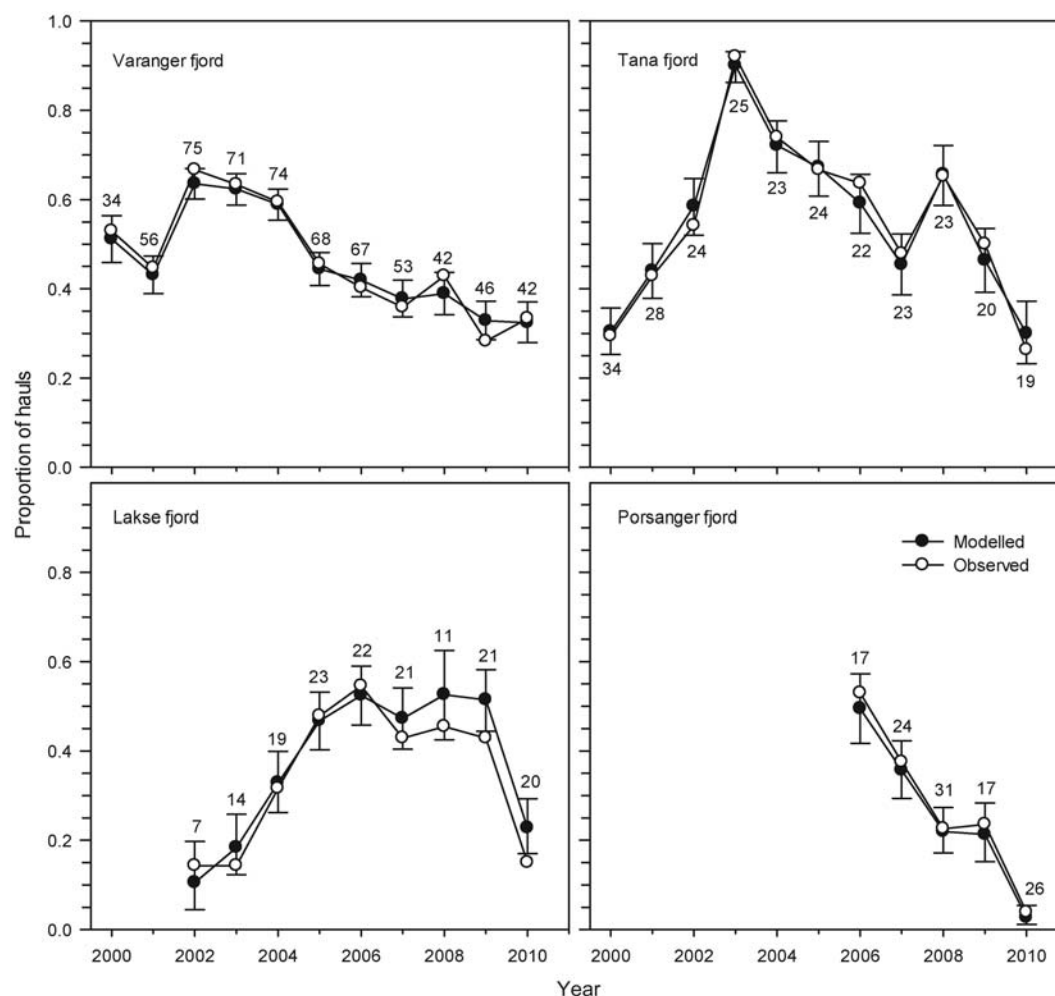


Figure 5. King Crabs in northern Norway: observed and predicted proportion of non-empty hauls by year and fjord for depth class 2 (90–330 m). Error bars are inter-quartile range. Points are labelled with the number of stations trawled.

which obviously also bounded the posterior distribution (Figure 4). We inferred that the data from the non-empty hauls were even more skewed than could be exactly fitted by the generalized gamma distribution with this limit on λ , and fitted lognormal distributions as a check on the results.

The depth range 90–330 m dominated the survey area with 69% of its extent; 28% of the extent was up to 90 m deep and only 3% deeper than 330 m (Table 1). 81% of all the stations, and 89% of the stations where crabs were caught, were in 90–330 m depths, which therefore dominated the estimation of year–area means both for presence–absence and for density when present, as they also dominated the consequent calculation of the total numbers of crabs in the survey area. The year–area and depth effects were fitted without weighting the data, which implied assuming that the density of stations was nearly the same in all sampling units. The relatively simple presence–absence model fitted very well to the observations for depth class 2, less well to depth classes 1 and 3.

The proportion of occupied sites (non-empty hauls) in the three depth strata follow a consistent pattern over time. Between fjords, however, the time trends differ (Figure 5). In Varanger and Tana fjords, the proportion of occupied sites peaked in 2002 and 2003 respectively, and decreased thereafter (Figure 5). In

Lakse fjord it increased until 2006, remained stable until 2009, and then decreased; in Porsanger fjord it decreased after 2006. In general, the proportion of occupied sites are lowest in depths <90 m – about one-third of the probabilities found in the two other depth strata (90–330 m and >330 m).

The more complex model for density in occupied areas fitted common values for the two shape parameters to all the sampling units, the scale parameter being fitted with two crossed sets of main effects. This more complex model was more difficult to evaluate. The density data had moderate variance—the overall CV was 1.77 and the skewness 7.04; within sampling units the calculated CVs ranged up to 2.28, but skewness only to 4.63; for depth class 2, skewness and CV of observed density had a correlation of 85.6%. The median estimate of the CV of the generalized gamma distribution fitted to the data was 1.03 with relative inter-quartile range (IR) 6%, but the distribution was highly skewed—using an approximation appropriate to a lognormal distribution, the median estimate of the skewness of the fitted distribution was 4.4 (IR = 9%). An ordered cumulation of probabilities that observed station densities would exceed the corresponding fitted values showed a deficit of very small values of observed density, probably due to a one-crab minimum catch, a surplus of moderately small values, and at the high end a surplus of very high values

(Figure 6). The model generally predicted smaller mean densities than those observed in sampling units with high CVs of density, and vice versa (Figure 7).

The lognormal model with sampling-unit CV related to mean density fitted better than the generalized gamma with assumed

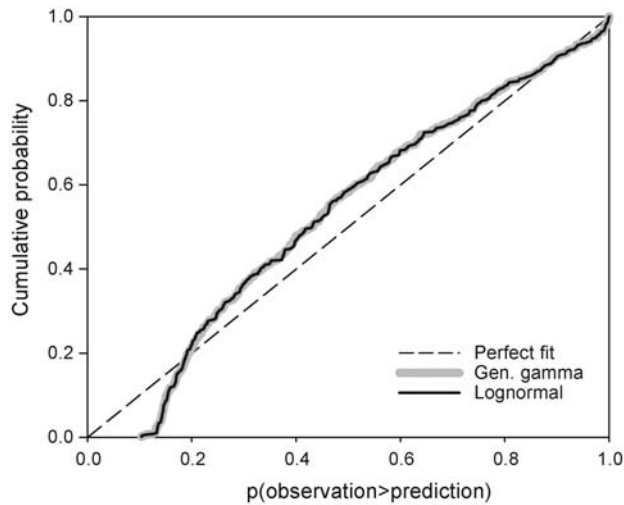


Figure 6. King Crabs in northern Norway: cumulative probability density distributions of probabilities that observed densities in non-empty hauls exceed corresponding values predicted by the generalised gamma – and lognormal models.

uniform CV (Table 2, Figure 7). (A simpler version with a common ratio of CV to mean had the same deviance as the version with depth-specific coefficients, and two fewer parameters, so was preferred.) It had a much smaller mean deviance and also estimated that it was effectively fitting four fewer parameters. The drawback to such a model is that the estimated CVs exercise a reverse influence on the sampling-unit means through the fitted relationship. In other words, fitting a model in which CV of density is related to its mean causes the estimated mean for a sampling unit to depend not only on the central tendency of the values observed, but also on their scatter; and *vice versa* the estimate of variance depends not only on the scatter of the observations, but also on their central tendency.

In spite of the differences in model diagnostics the generalized gamma and the lognormal models gave similar results. The overall mean densities of crabs ≥ 137 mm CL in the four fjords estimated by the model ranged from 14 km^{-2} to about 330 km^{-2} (Figure 8). The median CV of these annual estimates was 33%, with a range from 15% to 181%. The density estimates for the Varanger fjord were the most precise, with CVs from 15 to 31%, as they were based on more observations than those for the other fjords.

The estimated density indices for the four fjords have different time trends, although in Varanger, Tana and Lakse fjords a peak was apparently reached in 2003 or 2004. In Varanger fjord, densities increased from around 160 km^{-2} in 2001 to a plateau near 230 km^{-2} in 2002–2004 (Figure 8). Since then densities have progressively decreased to reach about 100 km^{-2} in 2010.

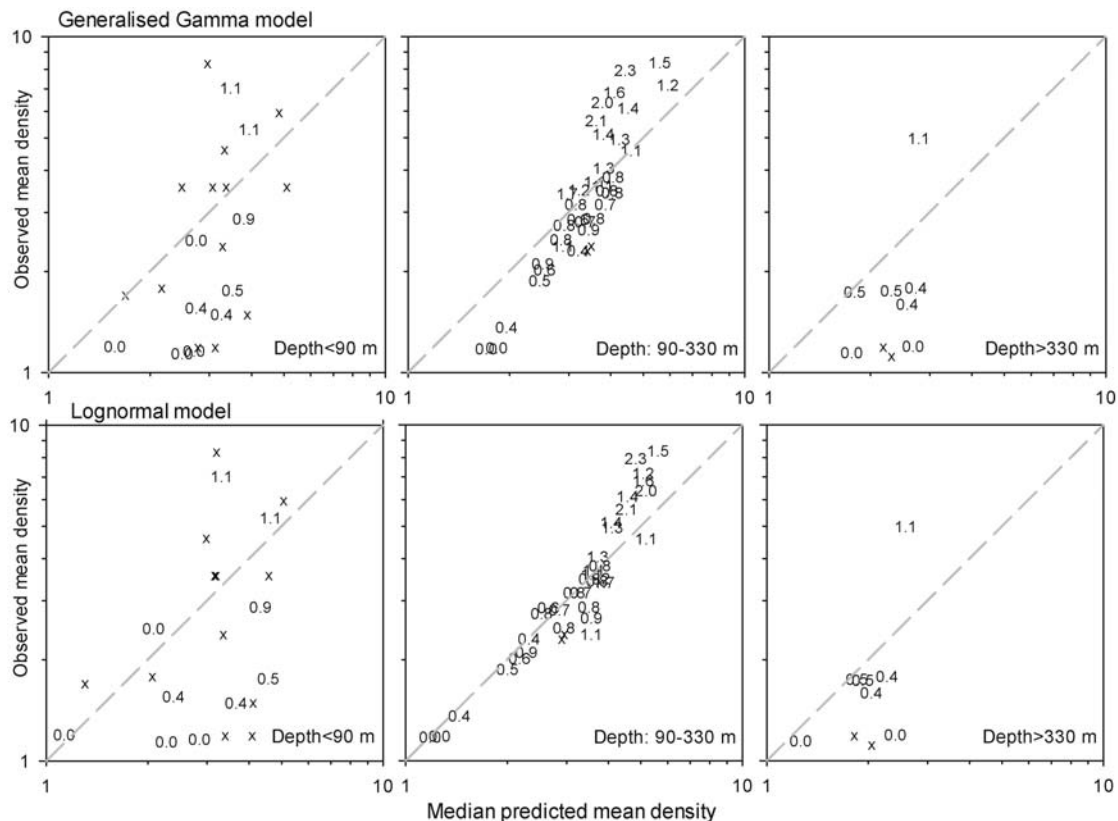


Figure 7. King Crabs in northern Norway: observed and predicted mean density in non-empty hauls by depth strata for the generalised gamma model (top row) and the lognormal model with CV proportional to the mean (bottom row). Points are labelled with the CV of the observed station density, points without CV (only one haul) are labelled with “x”.

In Tana fjord densities increased rapidly after 2000 to peak in 2003 near 300 km^{-2} , but then progressively declined to about 110 km^{-2} in 2007. The estimates for 2009 and 2010 are lower again, but there was a higher, although uncertain, estimate in 2008.

No crabs were detected before 2002 in Lakse fjord, but mean densities increased to about 130 km^{-2} by 2004 and stayed

relatively stable near that level through 2009, to decrease markedly in 2010. However, all estimates had large uncertainties. In Porsanger fjord large crabs were first recorded in 2006 with a density near 150 km^{-2} . The estimates of the following three years have been consecutively lower and Porsanger fjord currently has the lowest densities of the four.

While the crab stock gradually expanded its distribution westward the total number also increased (Figure 9). However after 2003 stock size has been progressively decreasing, by 2010 to about 36% of its 2003 maximum. Median annual estimates from the generalized gamma and lognormal models ranged from 0.4 to 1.1 million, with CVs of 12–24%.

Table 2. King crabs in Northern Norway: fits of distributions to a two-way GLM–delta model, sampling unit (i.e. year and fjord) crossed with depth.

	Mean Deviance	pD	DIC
Gamma, uniform CV	3 975.70	76.937	4 052.64
Gamma, with $\text{CV} \propto \text{mean}$	3 858.71	74.733	3 933.44
Generalized gamma, uniform CV ^a	3 758.48	77.501	3 835.98
Lognormal, uniform CV	3 752.67	77.316	3 829.99
Lognormal with $\text{CV} \propto \text{mean}$	3 662.00	73.758	3 735.76

^aThe generalized gamma was constrained by computational difficulties and was not a completely free fit.

Discussion

The surveys for king crab off Northern Norway use a modified Agassiz trawl as the primary sampling gear. The advantages of this gear over an ordinary (otter) trawl are that it is easy to operate, needs only small areas of trawlable bottom and is thus better suited to the rough bottom conditions found in the king

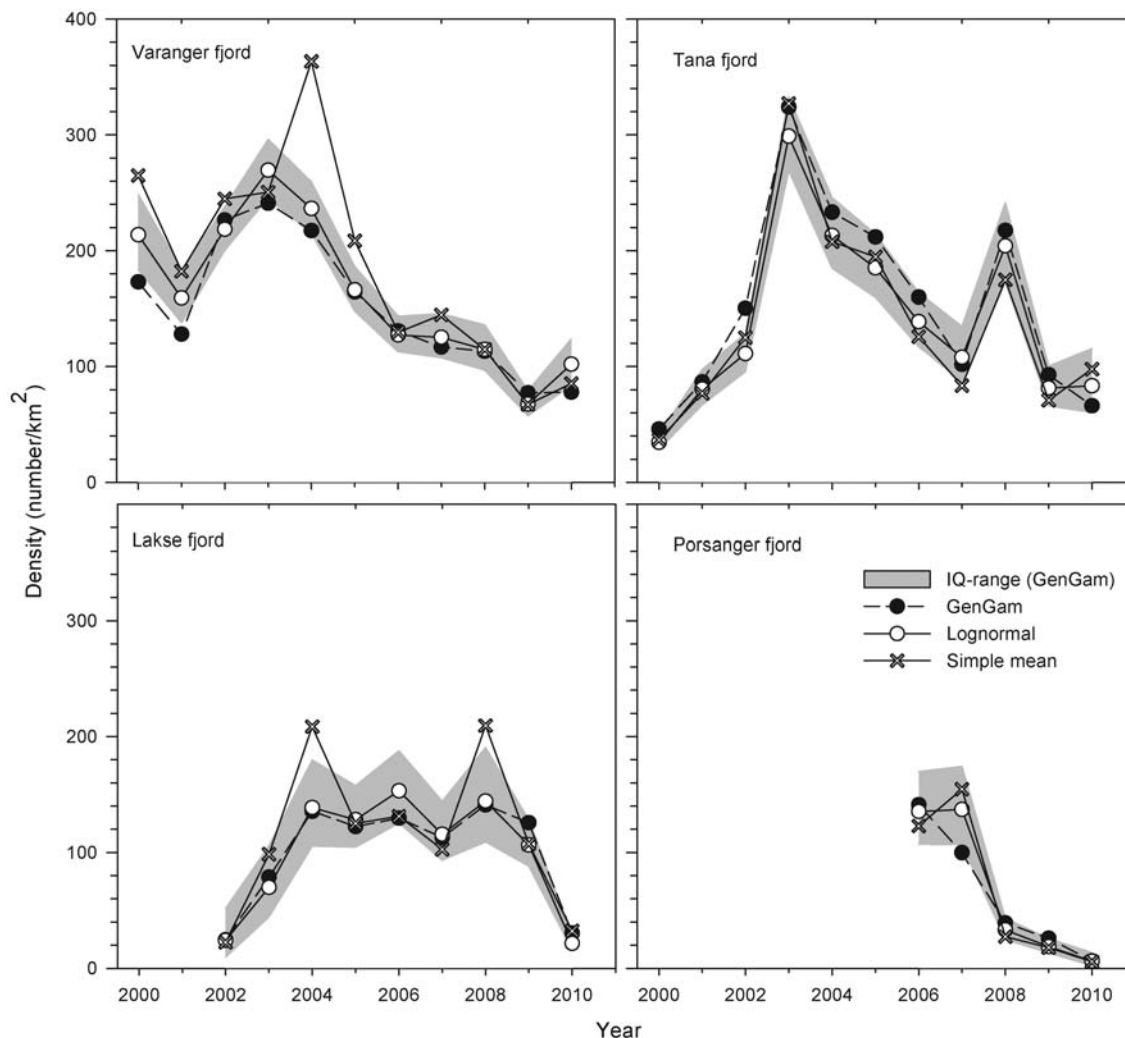


Figure 8. King Crabs in northern Norway: overall estimates of densities of crabs $> 137 \text{ mm}$ CL in each of four fjords, 2000–2010 from a) the generalised gamma model, b) the lognormal model with CV related to mean and c) calculated from simple means in sampling units. The greyed area shows the inter-quartile range as estimated from the lognormal model. The dots indicate medians of the estimate of the means.

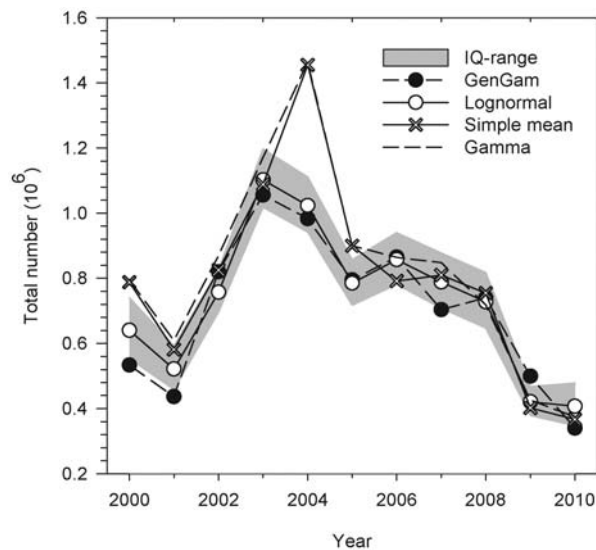


Figure 9. King Crabs in northern Norway: estimated total number of king crabs >137mm CL in four fjords, 2000–2010, from a) the generalised gamma model, b) the lognormal model with CV related to mean, c) calculated from simple means in sampling units and d) a model using gamma distribution with uniform CV. The greyed area shows the inter-quartile range as estimated from the lognormal model. The dots indicate medians of the estimate of the means.

crab habitats off Northern Norway. The use of the Agassiz trawl facilitates almost complete survey coverage of the king crab distribution and thereby reduces unknown biases on stock-size estimates. Such bias could originate from incomplete coverage, e.g. imposed by use of less adapted trawling gears (otter trawls), and crab migration between trawlable and non-trawlable areas. The Agassiz trawl is also effective in catching organisms in, on and very close to the bottom. Underwater video has shown that on soft bottoms the trawl catches benthos inhabiting the upper 5–10 cm of the sediment, and this is confirmed by finding in-benthic bivalve molluscs such as *Clinocardium ciliatum* and *Arctica islandica*, as well as several species of sedentary polychaete in the catches.

The data collected by these surveys are distributed with a long positive tail and include zeros in excess of what may be fitted by standard distributions. Such data are typical in fish stock surveys and may be modelled in a coherent way as a mix of two (or more) distributions. For the data from trawl surveys of king crab in waters off Northern Norway this was accomplished using a mixture of two general linear models—one for occurrence (zeros and non-zeros) and a second for non-zero catches—and Bayesian inference.

In previous analyses of this kind, the distribution of the positive data has often been taken as either lognormal or gamma (Syrjala, 2000; Dick, 2004). Lognormal and gamma models (or models assuming other distributions) give different results as also seen in this study (Figure 9) and there has been some debate over which of the two to prefer (Myers and Pepin, 1990; Pennington, 1996; Syrjala, 2000; Dick, 2004). The generalized gamma distribution has an extra parameter to regulate skewness and allows a model to fit within a spectrum of distributions that includes both the lognormal and the gamma. The gamma distribution has skewness equal to twice the CV; the lognormal has skewness equal to

three times the CV plus its cube. The generalized gamma can fit within this gap in the spectrum of skewness. In these analyses, the regulating parameter was given an uninformative prior, and the model was allowed to find the best fitting distribution. We find this approach useful when analysing such data: In the cases where neither the lognormal or gamma (or other) is a clear best fit, the generalized gamma distribution might provide an intermediate solution; if on the other hand the results of a first run using the generalized gamma distribution point to a specific standard two-parameter distribution as the better one it provides direction and statistical basis for simplifying the model.

In the present case, the parameter estimates of the generalized gamma model resulted in a distribution that tended strongly towards the lognormal, and there was clear evidence that the data fitted better to lognormal distributions than to e.g. gamma distributions (Table 2). The difference between lognormal and generalized gamma was negligible (Table 2). For practical reasons the simpler model with lognormal distribution was therefore preferred. It was however further modified to accommodate the CV being proportional to the mean (Figure 3). This was not possible with the generalized gamma model. The lognormal model fitted with CV proportional to mean density fitted better than either of the other models (Figure 6), with a smaller deviance (Table 2). However, for these models the resulting estimates of crab densities were quite similar (Figures 8 and 9).

The method used here is not error-free. The conceptual model underlying the method divides the habitat into “unoccupied” areas—in which the density, and therefore the catch, of crabs is everywhere zero—and “occupied” areas in which the catch has a certain distribution, modelled as continuous. However, the catch of crabs is in fact integer and forms a discrete distribution, many catches comprised only one crab, and it is statistically certain that some catches in “occupied” areas will have been zero. These catches are indistinguishable from, and were therefore included with, catches in “unoccupied” areas. This will especially be so in habitat areas occupied at low mean density. As a result, the fraction of habitat that is “occupied” will be underestimated, but at the same time the mean density in the “occupied” habitat will be overestimated because all the zero catches will have been sent off to be included with those from the “unoccupied” areas. These errors approximately cancel out: a rigorous analysis should estimate

$$p_{occ} = \frac{N_+ + N_0}{N},$$

$$\bar{O} = \frac{\sum_+ C + \sum_0 C}{N_+ + N_0}$$

and

$$\bar{D} = p_{occ} \cdot \bar{O} = \frac{\sum C}{N}$$

where p_{occ} is the proportion of habitat that is occupied, N_+ is the number of non-zero catches, N_0 is the number of zero catches in occupied habitat, with $\sum C$ being the corresponding sums of catches, N is the total number of catches, \bar{O} is the estimated mean density in occupied habitat and \bar{D} the overall mean

density in the sampling unit. The analysis described here in fact estimates

$$p_{occ} = \frac{N_+}{N}$$

and

$$\bar{O} = \frac{\sum C}{N_+}$$

but still arrives approximately at

$$\bar{D} = p_{occ} \cdot \bar{O} = \frac{\sum C}{N}.$$

The error incurred lies in fitting a distribution to truncated data without taking account of the truncation, so estimates of the parameters of the distribution, and even statements about which distribution fits best, are not necessarily reliable.

The distribution of data

For given mean and variance, the lognormal is always skewed at least half as much again as the gamma distribution, and often much more. For this reason the generalized gamma, capable of generating a spectrum of distributions of intermediate skewness, can be useful in fitting to skewed data. However, the expressions involved in fitting the distribution—either $\ln\left(\Gamma\left(\kappa + \frac{1}{\beta}\right)\right) - \ln(\Gamma(\kappa))$ or $\sqrt[\beta]{\kappa}$ —can present computational problems for large κ and small β , and distributions close to lognormal are computationally difficult of access. The parameter λ affects the skewness, but the relationship is not simple, and skewness is also affected by sigma. The limit as λ tends to 0 (large κ and small β) is a lognormal distribution, with skewness equal to $CV^3 + 3 \times CV$, and $\lambda = \sigma$ (i.e. $\beta = 1$) generates an ordinary gamma distribution, with skewness equal to twice the CV. If $\lambda\sigma \geq 1$ (i.e. $\beta\kappa \leq 1$) the distribution is J-shaped. J-shaped distributions of non-zero density were unwelcome in this analysis and λ was therefore limited to values less than $1/\sigma$. However, it was not considered necessary to restrict distributions to those between gamma and lognormal, so λ would have been allowed to exceed sigma (i.e. $\beta > 1$) although it did not want to. The negative binomial, with “zero inflation”, has been considered a candidate for fitting to data from biological surveys (e.g. Martin *et al.*, 2005; Minami *et al.*, 2007; Joseph *et al.*, 2009), but it has two limitations. Firstly, it has only two parameters and so lacks the freedom to fit to three properties—mean, variance and skewness—of survey data distributions; and second, its skewness, ranging from CV to $2 \times CV$, is even less than that of the gamma distribution and cannot reach the values often encountered in survey data.

Results were also computed using the simplest method: calculating the mean untransformed density for all stations in a sampling unit, and estimating the mean density in each fjord, and the total number of crabs in the survey area, by ordinary arithmetic. Error variances and standard errors of density and of the total numbers of crabs were also calculated by standard statistical methods. There were large differences in some years between results of these calculations and those from the more complex model, especially in 2004 (Figures 7 and 8).

There are several reasons for these differences to occur. First, the estimation methods were different. The simple-means method corresponds to maximum likelihood under an assumption that data are normally distributed, conditions that clearly do not hold in these data, whereas the model that we fitted by Bayesian methods assumes a skewed distribution for the data. The simple-means method is sensitive to outlying values, while the fitted model expects positively skewed data and is more robust to positive outliers. Second, the CV model was different: the simple-means approach assumes nothing about any relationship between variance and mean. Contrariwise, the fitted model assumes a constant CV, or SD proportional to mean, and will therefore attempt to adjust its estimates of cell mean values to fit with cell SDs; and *vice versa*. Third, the model assumptions were different. Calculating a simple mean assumes that each station, whether or not it caught crabs, provided an unbiased estimate of the mean density in that sampling unit alone, and that all stations in the sampling unit had the same expectation of squared density deviation from the mean. The simple-mean estimates are unbiased and assumption-free, but these are their only virtues: they are not necessarily in the neighbourhood of the most likely values. The modelling method, with its two-way analysis, assumed that sampling units shared common depth effects and common year-area effects, both as regards the proportions of the sampling units that were occupied and the density in occupied areas.

Thus, for these various reasons, it could be expected that different analysis methods would give different results without necessarily implying that one or the other was “wrong”, although between the two modelling approaches, which had a common measure of fit, the lognormal with variable CV appeared to fit better. The simple-means method analysed the data for each sampling unit independently of the analyses for the others. The modelling method, by contrast, fitted common depth effects over all years and areas, and this could be expected to reduce the year-to-year variation in estimates for the modelling method, while the simple-means method was allowed to respond more erratically to each year’s individual results. Over the 11-year series of surveys, the log-normal and generalized gamma models estimated lower average total numbers than the simple-means method, but medians of distributions estimated by Bayesian methods are not unbiased estimates of corresponding means.

Notwithstanding the above, the trajectories of estimated total numbers in the survey area were similar, except for a large positive excursion by the simple-means estimate in 2004 (Figures 7 and 8). In that year, two of the three largest catches in the entire survey series were made in depth class 2 in Varanger fjord, which accounts for 33% of the survey area, or without Porsanger fjord, which was not surveyed in 2004, 42%.

However, we also tried fitting an ordinary gamma distribution, with uniform CV, in the two-way model. The trajectory of total numbers was strikingly close to that given by the simple-means method (Figure 9). This appears to indicate that the divergence in the 2004 estimate is not only due to the difference between a structured model and a one-sampling-unit-at-a-time analysis method. Instead, it appears that models based on distributions that are inconsistent with the way the data are distributed, with inadequate ability to recognize and fit to its high positive skewness, give markedly different results from others that are more consistent with the highly skewed data. There is also a difference between estimates that intend to be unbiased and others that have other objectives.

The estimated crab stock

The estimated densities of the large crabs (>137 mm CL) in the four fjords vary as the king crab expands its distribution westward from the points of release in the eastern Barents Sea (Figure 1). When the survey series started in 2000 the king crab was already well established in Varanger fjord and the year-to-year variation shown by the survey (Figure 8) can probably be attributed to natural fluctuation and to the effect of the fishery. In Tana fjord the first ripe female crabs were reported in 1995, and therefore the 2000–2003 increase in density is probably the completion of the initial population growth due to colonization. The fishery was not fully established here until 2001. Further east, in Lake fjord, densities were zero before 2002 after which the crab began to colonize this fjord as well. In 2003 a fishery had already begun. With a further delay, this course of events was repeated in Porsanger fjord.

The estimated large crab densities ranged up to over 300 km^{-2} , but most values were between 100 km^{-2} and about 250 km^{-2} . The catchability of crab by the sled gear that was used is believed to be close to 100%. However, as is usually the case for fisheries surveys, hard quantitative estimates of catchability are not available, and obtaining such estimates, if the catchability really is close to 100%, would be difficult. The survey results are considered reliable as regards relative values over time or in different areas.

The overall CV of the survey has been as low as 13%, but in recent years has hovered nearer to 17–20%. This is perhaps because densities are now below their peak values and it is more difficult to get small CVs when resources become scarcer, but the taking of fewer hauls in recent years probably also contributes to the increased CV. Even so, CVs in these ranges compare well with those of most resource surveys.

The methods described provide estimates of the full probability density distribution of the annual stock size. Thus, it is straightforward to use the results directly in management evaluations through risk analyses to help fisheries management make decisions.

Acknowledgements

We would like to thank the people conducting the surveys: crew, technicians and scientists.

References

- Berenboim, B. I., and Pinchukov, M. A. 2005. Status and measures of red king crab stock management in Russian waters of the Barents Sea. In *Investigations on the Kamtschatka Crab (Paralithodes camtschaticus) in the Barents Sea during 2002–2004*. Ed. by Sundet, J. H. Report to the 34th Session of the Joint Norwegian–Russian Fishery Commission, Kaliningrad, October 2005. 41 pp.
- Congdon, P. 2001. *Bayesian Statistical Modelling*. Wiley, London. 556 pp.
- Dick, E. J. 2004. Beyond 'lognormal versus gamma': discrimination among error distributions for generalized linear models. *Fisheries Research*, 70: 351–366.
- Fletcher, D., MacKenzie, D., and Villouta, E. 2005. Modelling skewed data with many zeros: A simple approach combining ordinary and logistic regression. *Environmental and Ecological Statistics*, 12: 45–54.
- Hjelset, A. M., Sundet, J. H., and Nilssen, E. M. 2009. Size at sexual maturity in the female red king crab (*Paralithodes camtschaticus*) in a newly settled population in the Barents Sea, Norway. *Journal of Northwest Atlantic Fishery Science*, 41: 173–182.
- Joseph, L. N., Elkin, C., Martin, T. G., and Possingham, H. P. 2009. Modelling abundance using N-mixture models: the importance of considering ecological mechanisms. *Ecological Applications*, 19: 631–642.
- Martin, T. G., Wintle, B. A., Rhodes, J. R., Kuhnert, P. M., Field, S. A., Low-Choy, S. J., Tyre, A. J., et al. 2005. Zero tolerance ecology: improving ecological inference by modelling the source of zero observations. *Ecology Letters*, 8: 1235–1246.
- Maunder, M. N., and Punt, A. E. 2004. Standardizing catch and effort data: a review of recent approaches. *Fisheries Research*, 70: 141–159.
- McArdle, B. H., and Anderson, M. J. 2004. Variance heterogeneity, transformations, and models of species abundance: a cautionary tale. *Canadian Journal of Fisheries and Aquatic Sciences*, 61: 1294–1302.
- Minami, M., Lennert-Cody, C. E., Gao, W., and Roman-Verdesoto, M. 2007. Modelling shark bycatch: the zero-inflated negative binomial regression model with smoothing. *Fisheries Research*, 84: 210–221.
- Myers, R. A., and Pepin, P. 1990. The robustness of lognormal based estimators of abundance. *Biometrics*, 46: 1185–1192.
- Otto, R. S. 1986. Management and Assessment of Eastern Bering Sea king crab. In *North Pacific Workshop on Stock Assessment and Management of Invertebrates*. Ed. By Jamieson, G. S., and Bourne, N. Canadian Special Publications of Fisheries and Aquatic Sciences 92.
- Pennington, M. 1983. Efficient estimators of abundance, for fish and plankton surveys. *Biometrics*, 39: 281–286.
- Pennington, M. 1996. Estimating the mean and variance from highly skewed marine data. *Fishery Bulletin*, 94: 498–505.
- Prentice, R. L. 1974. A log gamma model and its maximum likelihood estimation. *Biometrika*, 61: 539–544.
- Shono, H. 2008. Application of the Tweedie distribution to zero-catch data in CPUE analysis. *Fisheries Research*, 93: 154–162.
- Sokal, R. R., and Rohlf, F. J. 1995. *Biometry: the Principles and Practice of Statistics in Biological Research*, 3rd edn. W. H. Freeman and Co., New York. 887 pp.
- Spiegelhalter, D. J., Thomas, A., Best, N., and Lunn, D. 2004. WinBUGS User Manual, version 1.4. January 2003. MRC Biostatistics Unit, Institute of Public Health, Cambridge, England.
- Stacy, E. W. 1962. A generalisation of the gamma distribution. *The Annals of Mathematical Statistics*, 33: 1187–1192.
- Stacy, E. W., and Mihram, G. A. 1965. Parameter estimation for a generalised gamma distribution. *Technometrics*, 7: 349–358.
- Stefánsson, G. 1996. Analysis of groundfish survey abundance data: combining the GLM and delta approaches. *ICES Journal of Marine Science*, 53: 577–588.
- Syrjala, S. E. 2000. Critique on the use of the delta distribution for the analysis of trawl survey data. *ICES Journal of Marine Science*, 57: 831–842.
- Zheng, J., Murphy, M. C., and Kruse, G. H. 1997. Analysis of harvest strategies for red king crab, *Paralithodes camtschaticus*, in Bristol Bay, Alaska. *Canadian Journal of Fisheries and Aquatic Sciences*, 54: 1121–1134.
- Zuur, A. F., Ieno, E. N., Walker, N. J., Saveliev, A. A., and Smith, G. M. 2009. Zero-truncated and zero-inflated models for count data. In *Mixed Effects Models and Extensions in Ecology with R*, pp. 261–293. Springer, New York. 574 pp.

Handling editor: Sarah Kraak