



Model Selection via Focused Information Criteria for Complex Data in Ecology and Evolution

Gerda Claeskens¹, Céline Cunen² and Nils Lid Hjort^{2*}

¹ ORSTAT and Leuven Statistics Research Centre, KU Leuven, Leuven, Belgium, ² Department of Mathematics, University of Oslo, Oslo, Norway

Datasets encountered when examining deeper issues in ecology and evolution are often complex. This calls for careful strategies for both model building, model selection, and model averaging. Our paper aims at motivating, exhibiting, and further developing focused model selection criteria. In contexts involving precisely formulated interest parameters, these versions of FIC, the focused information criterion, typically lead to better final precision for the most salient estimates, confidence intervals, etc. as compared to estimators obtained from other selection methods. Our methods are illustrated with real case studies in ecology; one related to bird species abundance and another to the decline in body condition for the Antarctic minke whale.

OPEN ACCESS

Edited by:

Jose Miguel Ponciano,
University of Florida, United States

Reviewed by:

Subhash Ramkrishna Lele,
University of Alberta, Canada
Christopher Jackson,
University of Cambridge,
United Kingdom

*Correspondence:

Nils Lid Hjort
nils@math.uio.no

Specialty section:

This article was submitted to
Environmental Informatics,
a section of the journal
Frontiers in Ecology and Evolution

Received: 27 February 2019

Accepted: 17 October 2019

Published: 14 November 2019

Citation:

Claeskens G, Cunen C and Hjort NL
(2019) Model Selection via Focused
Information Criteria for Complex Data
in Ecology and Evolution.
Front. Ecol. Evol. 7:415.
doi: 10.3389/fevo.2019.00415

Keywords: bird species abundance, ecology, evolution, FIC and AFIC, focused model selection, linear mixed effects, minke whales

1. INTRODUCTION

Only rarely will initial modeling efforts lead to “one and only one model” for the data at hand. This simple empirical statement applies in particular to situations with complex data for complicated and not-yet-understood mechanisms underlying the phenomena being studied, in ecology and evolution, as well as other sciences. Thus, methods for model comparison, model selection, and model averaging are called for. Not surprisingly there must be several such methods, since the question “what is a good model for my data?” cannot be expected to have a simple and clear-cut answer.

There are indeed several model selection schemes in the statistics literature, with the more famous ones being the AIC (the Akaike Information Criterion) and the BIC (the Bayesian Information Criterion; see Claeskens and Hjort, 2008b) for a general overview. The AIC and BIC are able to compare and rank competing models for a given dataset, as long as they are all parametric. These and yet other methods work in an “overall modus,” in appropriate senses comparing overall fit with overall complexity, but they do not take on board *the intended use of the fitted models*. This is where FIC (the Focused Information Criterion) comes in, along with certain relatives. The FIC aims at giving the most relevant model comparison and ranking, and hence also pointing to the best model, for *the given purpose*. What this given purpose is depends on the scientific context. Indeed, two research teams might ask different focused questions, with the same data and the same list of candidate models, and we judge it not to be a contradiction in terms that three focused questions might have three different best models.

The present article gives an account of FIC and its relatives, including also certain extensions of previously published methods. We do have models for ecology and evolution in mind, though it is clear that the view is broader: we wish to find good statistical models for complex data, and

can do so, once crucial and context driven questions are translated to *focus parameters*. Our paper’s contribution is 2-fold. (i) We aim at introducing the FIC methodology to researchers in ecology and evolution. We have therefore strived to include relevant examples, along with some R code. We also discuss various topics of interest to applied researchers, particularly in section 5. In this partly tutorial spirit, various technical details have been placed in the **Appendix**. (ii) Our article also serves as an outlet for a somewhat new FIC framework, termed the “fixed wide model framework,” different from the “local asymptotics framework” used in the majority of previous publications. Details are in section 3, with material not been presented in this general form before. In particular, the extension of this framework to generalized linear models is novel.

To help fix ideas and some basic notation, we start with a concrete application. We use the dataset from Hand et al. (1994) regarding counts of the number of bird species on fourteen areas, vegetation islands, in the Andes mountains with páramo vegetation. In addition to the number of bird species y , there are four covariates recorded for each such vegetation island: x_1 , the area of the vegetation island in thousands of square kilometers; x_2 , the elevation in thousands of meters; x_3 , the distance between the area and Ecuador in kilometers; and x_4 , the distance from the nearest island in kilometers.

y	x_1	x_2	x_3	x_4
36	0.33	1.26	36	14
30	0.50	1.17	234	13
37	2.03	1.06	543	83
35	0.99	1.90	551	23
11	0.03	0.46	773	45
21	2.17	2.00	801	14
11	0.22	0.70	950	14
13	0.14	0.74	958	5
17	0.05	0.61	995	29
13	0.07	0.66	1065	55
29	1.80	1.50	1167	35
4	0.17	0.75	1182	75
18	0.61	2.28	1238	75
15	0.07	0.55	1380	35

We model the number of bird species Y by a Poisson distribution with mean $\exp(x^t\beta)$, where x in the widest model consists of the constant 1 (modeling the intercept), all four covariates x_1, \dots, x_4 as main effects, and all six pairwise interactions between these main effects. This amounts to a total of 11 parameters $\beta_0, \dots, \beta_{10}$. We wish to include the intercept parameter β_0 in all candidate models, and hence take it as a “protected parameter,” whereas the other parameters are “open,” and can be pushed in and out of candidate models. For this application, all submodels of the largest 11-parameter model are considered, with the further restriction that interactions between two covariates can be included only if the two main effects are present. This results in a total of 113 models.

The main distinction between FIC and various other information criteria is the presence of a *focus*. This is a quantity of interest that depends on the model parameters and is estimable from the data. The generic notation for the focus used in our paper is μ . Its dependence on the model parameters might be indicated by writing $\mu(\beta)$.

In the bird species study, our first focus concerns one of the vegetation islands, Chiles. This area is the one among the fourteen that is closest to Ecuador, and has covariate values $x_1 = 0.33, x_2 = 1.26, x_3 = 36, x_4 = 14$. We wish to select a model that best estimates the expected number of bird species for this island,

that is, $\mu(\beta) = \exp(x^t\beta)$ for the given covariate values for Chiles. In our model search problem there are 113 models and hence 113 estimators for μ . Each such estimator, say $\hat{\mu}_M$ for a candidate model M , comes with its own bias and variance, say b_M and τ_M^2 . Thus, for each candidate model there is a corresponding mean squared error (mse)

$$\text{mse}_M = \tau_M^2 + b_M^2. \tag{1}$$

The basic idea of the FIC is to estimate these mse values from the data, for the wide as well as for each candidate model, i.e., to construct

$$\text{FIC}_M = \widehat{\text{mse}}_M = \widehat{\tau}_M^2 + \widehat{\text{bsq}}_M, \tag{2}$$

with the second term indicating estimation of the squared bias $\text{bsq}_M = b_M^2$. In the end one selects the model with the smallest estimated mse.

For the bird species application, we use FIC for finding the best model to estimate the expected number of bird species for Chiles. We use the R package `fic` with the following lines of R code, where we fit the wide model, specify the focus function, the covariate value in which to evaluate this focus, and the specific models that we wish to search through. In this example we restrict the built-in all subsets specification to only using models that obey the hierarchy principle (so out of the $2^{10} = 1024$ potential submodels, only the 113 pointed to above are included).

```
library(fic)
wide.birds = glm(y~.^2, data=birds,
  family=poisson)
focus1 = function(par, X) exp(X %*% par)
inds0 = c(1,rep(0,10)) # only the intercept
is in the narrow model
A = all_inds(wide.birds, inds0) # use all
subsets of the wide model
#exclude models with interactions that do
not have both main effects:
inds <- with(A,A[(A[,2]==0 & (A[,6]==
1|A[,7]==1|A[,8]==1) |
  A[,3]==0 & (A[,6]==1|A[,9]==1
  |A[,10]==1) |
  A[,4]==0 & (A[,7]==1|A[,9]==1
  |A[,11]==1) |
  A[,5]==0 & (A[,8]==1|A[,10]==1|
  A[,11]==1)], )
# specify the X used to evaluate the focus
function:
XChiles=model.matrix(wide.birds)[1, ]
fic(wide=wide.birds, inds=inds, inds0=inds0,
  focus=focus1, X=XChiles)
```

For each of the 113 models we get via the output values of the focus estimate, the estimated bias, standard error, and actually two versions of the FIC of (2), corresponding to two related but different ways of estimating the b_M^2 part (for details, see section 2). For FIC tables and FIC plots we prefer working with the square-root of the FIC, i.e., estimates of the root-mse (rmse) rather than

of the mse, as these are on the original scale of the focus and easier to interpret.

Table 1 is constructed from the output for a selection of models, including the narrow model (1) which has a relatively large (in absolute value) bias estimate of -19.035 , a relatively small standard error of 2.247 and a focus estimate of 20.71 ; the wide model (113) with zero as the bias estimate though with a large standard error of 6.051 . This is a typical output: the wide model contains 11 parameters to estimate which causes the standard error to be large, the narrow model only contains the intercept resulting in a small standard error. For the bias estimate the scenario is reversed: the wide model has the smallest bias, while the narrow model has a larger bias. The balancing act of the FIC via the mean squared error finds a compromise. The selected model (5) results in the smallest value of the square root of the estimated mean squared error (rmse). Its indicator sequence 10010,000000, with a one for β_0 and β_3 , and zeroes for the interactions, points toward the selected focus $\mu(\beta) = \exp(\beta_0 + \beta_3 x_3)$ with corresponding estimated focus value 38.88 . Using the wide model would have resulted in a close 38.27 though with a larger estimated root mean squared error. The wide model only ranks at the 73rd place according to estimated rmse. Model (20) is selected by the Bayesian information criterion BIC, it consists of the intercept, all four main effects and the interaction between x_1 and x_2 . In the rmse ranking it comes at the 42nd place. Model (67) is the one selected by the Akaike information criterion, next to the intercept and all main effects it consists of the interactions $x_1 x_3, x_2 x_3, x_2 x_4$. This models ranks 32nd.

The second focus concerns the probability of having more than 30 bird species, $P(Y > 30 | x)$. Now we do not specify a particular island but use the average FIC (see section 2.2), with equal weights for the fourteen vegetation islands (non-equal weights can easily be worked with too).

```
focus2 = function(par, X) 1-ppois(30,
lambda=exp(X %*% par))
Xall = model.matrix(wide.birds)
fic2 = fic(wide=wide.birds,inds=inds,
inds0=inds0,focus=focus2,X=Xall)
AVE = fic2[fic2$vals=="ave",]
which.min(AVE$rmse.adj)
```

The AFIC selects the following form for the mean: $\exp(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_4 x_4 + \beta_7 x_1 x_4)$. The averaged focus estimate of the

probability of observing over 30 bird species in the selected model equals 15.73% , while the wide model's estimate is 21.83% , though with a substantial larger estimated mean squared error due to the estimation of 11 parameters instead of only 5 for the selected model. Of course, AIC and BIC ignore any information regarding the focus, and thus still recommend the very same models, model (67) for AIC, with estimate 21.15% , and model (20) for BIC, with estimate 21.59% . The AIC model ranks 16th, the BIC model is now at the third place.

Figure 1 displays for these two foci the root-FIC and root-AFIC values, as well as the estimated focus values, for all of the 113 models. The FIC or AFIC selected values, minimizing the respective criteria, are indicated in red, while the wide model's values are in blue.

Several traditional model selection criteria, such as the AIC and the BIC (see Claeskens and Hjort, 2008b, Chs. 2, 3) work in an overall modus, finding models that in a statistical sense are good on average, not taking on board the specific aims of a study. The FIC works explicitly with such specific aims, formalized via the focus parameters. Thus, FIC might find that one model works very well for covariates "in the middle," whereas another model could work rather better for covariates outside mainstream. Similarly, one model might work well for explaining means, and another for explaining variances. We stress that the FIC apparatus works for *any specified focus parameter*, and is not limited to e.g., regression coefficients and the customary selection of covariates from that perspective.

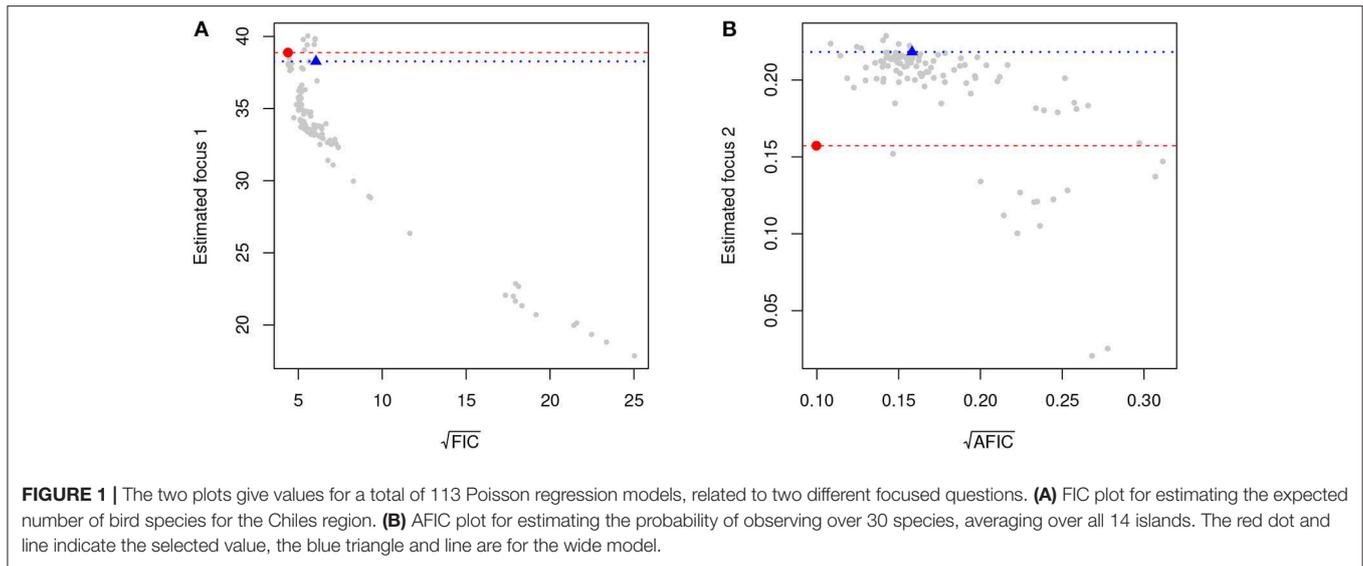
The generic FIC formula (2) cannot be immediately applied, as efforts are required to establish formulae for approximations to biases and variances, along with construction of estimators for these quantities. Thus, the FIC formula pans out differently in different situations, depending on the general framework, the complexity of models, and estimators of the focus parameters. A brief overview of general principles, leading to such approximations and estimators, is given in section 2. This also encompasses AFIC, ways of creating average-FIC scores in situations where more than one focus parameter is at stake.

In section 3 we provide the general FIC formulae in the so-called fixed wide model framework. The development of FIC formulae ingredients in a somewhat different framework, with local neighborhood models, is placed in **Appendix**. Generalized linear models are used as examples, encompassing linear regression, logistic and Poisson regression, etc. The more general

TABLE 1 | Bird species example.

Model	Coef. indicators	Focus	Bias	Se	\sqrt{FIC}	AIC	BIC
1	10000,000000	20.714	-19.035	2.247	19.167	143.26	143.90
5	<u>10010,000000</u>	<u>38.882</u>	0.000	4.383	<u>4.383</u>	112.65	113.93
20	11111,100000	33.718	-2.156	4.670	5.143	91.91	<u>95.74</u>
28	11101,001000	26.356	-11.0468	3.674	11.642	98.54	101.74
67	11111,010110	39.784	0.000	5.296	5.296	<u>91.44</u>	96.55
113	11111,111111	38.269	0.000	6.051	6.051	95.72	102.75

This table is constructed from output of the R function fic for six of the 113 models, together with the AIC and BIC values. FIC selection takes place via the square root of the estimated mean squared error of the focus estimator.



class of linear mixed models has proven important for various applications to ecology, and in section 3.3 FIC formulae are reached for such. In section 4 we use linear mixed effects models with FIC for analyzing the body conditions of minke whales in the Antarctic, where one focus parameter is the yearly decline in energy storage. A general but brief discussion is then offered in section 5. Here we touch on aspects of performance, along with a few concluding remarks, some of which point to future research.

2. FOCUSED INFORMATION CRITERIA

The application concerning birds on vegetation islands in the previous section was meant to provide intuition for the use of FIC for model selection. Here we give a more formal, but brief, overview of the FIC and AFIC schemes.

2.1. General FIC Scheme

Suppose we have defined a *wide model* which is assumed to be the true data-generating mechanism. Estimating the focus parameter using the wide model leads to $\hat{\mu}_{\text{wide}}$, which under broad regularity conditions will aim at μ_{true} , the unknown true value of the focus parameter. Estimation via fitting a candidate model M leads to $\hat{\mu}_M$, say, aiming for some least false parameter $\mu_{0,M}$, typically different from μ_{true} , due to modeling bias. The least false parameter in question relates to the best approximation candidate model M can manage to be, to the true model. There is therefore an inherent bias, say

$$b_M = \mu_{0,M} - \mu_{\text{true}},$$

associated with using M . We saw estimates of this bias in the birds application above, where small models could have larger biases.

The estimators will have certain variances. In most frameworks, involving independent or weakly dependent data, these tend to zero with speed $1/n$, in terms of growing sample size n . It is therefore convenient and informative to write these variances as $\tau_{\text{wide}}^2 = \sigma_{\text{wide}}^2/n$ and $\tau_M^2 = \sigma_M^2/n$, where the mathematics and approximation theorems associated with different frameworks typically yield expressions for or

approximations to the σ_{wide} and σ_M . The mse of the focus parameter estimators is the sum of the variance and the bias squared,

$$\text{mse}_{\text{wide}} = \sigma_{\text{wide}}^2/n + 0^2 \quad \text{and} \quad \text{mse}_M = \sigma_M^2/n + b_M^2. \quad (3)$$

These quantities are measures of the *risk*, in the statistical sense, associated with using each of the models for estimating μ . As explained in the introduction, the FIC scores of (2) are estimates of the mse of the focus parameter estimators, i.e., the $\hat{\mu}_M$, for a specific dataset, for each of the models under consideration. Equation (3) is also an informative reminder that with more data, variances get small, but biases remain. So using a model which is not fully correct can still yield sharper estimators, as long as the bias is moderate or small: $|b_M| < (\sigma_{\text{wide}}^2 - \sigma_M^2)^{1/2}/\sqrt{n}$. It is also clear that with steadily more data, steadily more sophisticated models can and indeed should be used. The FIC makes these ideas operative.

In various cases the variance terms σ_M^2/n are easier to estimate than the squared biases b_M^2 . A starting point for the latter is $\hat{b}_M = \hat{\mu}_M - \hat{\mu}_{\text{wide}}$, but the corresponding \hat{b}_M^2 will overshoot b_M^2 with about κ_M^2/n , which is the variance of b_M . With appropriately constructed estimators of the quantities σ_{wide} , σ_M , κ_M (with different recipes for different situations), this yields two natural ways of estimating the actual mse values:

$$\begin{aligned} \text{FIC}_{\text{wide}}^u &= \hat{\sigma}_{\text{wide}}^2/n + 0^2 & \text{and} & \quad \text{FIC}_M^u = \hat{\sigma}_M^2/n + \hat{b}_M^2 - \hat{\kappa}_M^2/n, \\ \text{FIC}_{\text{wide}} &= \hat{\sigma}_{\text{wide}}^2/n + 0^2 & \text{and} & \\ \text{FIC}_M &= \hat{\sigma}_M^2/n + \max(\hat{b}_M^2 - \hat{\kappa}_M^2/n, 0). \end{aligned} \quad (4)$$

The FIC^u scores are (approximately) unbiased estimates of the mse, since $\hat{b}_M^2 - \hat{\kappa}_M^2/n$ is (approximately) unbiased for b_M^2 , whereas the FIC scores are adjusted versions, by truncating any negative estimates of squared bias to zero, as we did in the first example. If the true bias in question is some distance away from zero, FIC_M^u will be equal to FIC_M . When faced with a specific application one should decide on one of these two FIC versions, and use the same choice for all models under consideration.

In order to turn the general scheme (4) into clear formulae, with consequent algorithms, we need expressions for or approximations to the population quantities σ_M , b_M , κ_M , followed by clear estimation strategies for these again. In most cases we need to rely on large-sample approximations. Arriving at clear formulae for σ_M etc. depends on the particularities of the wide model, the candidate models, and the focus parameter. We provide such FIC formulae, for *two different frameworks* or setups. *The first* involves local asymptotics, with candidate models being a local distance $O(1/\sqrt{n})$ away from the wide model. This derivation is placed in **Appendices A1** and **A2**. *The second* avoids such local asymptotics and works from a fixed wide model and a collection of candidate models (see section 3). It is not a contradiction in terms that these two frameworks lead to related but not identical FIC formulae, as different mathematical approximations are at work.

2.2. AFIC, the Averaged-Weighted Selection Scheme

The FIC apparatus above is tailored to one specific focus parameter at a time. In a regression context this applies e.g., to estimating the mean response function for one covariate vector at a time, say $\mu(\theta; x_0)$. Often there would be active interest in several parameters, however, as with such a $\mu(\theta; x_0)$ for all x_0 in a segment of covariates, or a probability $P(Y \geq y_0 | x_0)$ for a set of thresholds, as in the birds study.

Suppose in general that an ensemble of estimands is of interest, say $\mu(\theta; v)$ with $v \in V$, and that a measure of relative importance $dW(v)$ is assigned to these. There could be only a few such estimands under scrutiny, say μ_j for $j = 1, \dots, k$, along with weights of importance w_1, \dots, w_k . Estimation involving all higher quantiles, or all covariates within a certain region, however, would constitute examples where we need the more general $v \in V$ notation. Here we sketch the AFIC approach, for estimating the relevant integrated weighted risk.

For each focus parameter in the ensemble of estimands there is an associated mse or risk, $mse(v)$. The combined risk associated with using model M then becomes

$$r_n(M) = \int mse(v) dW(v) = \int \{\sigma_M(v)^2/n + b_M(v)^2\} dW(v),$$

with the appropriate $\sigma_M(v)$ and $b_M(v) = \mu_{0,M,n}(v) - \mu_{true}(v)$. An approximately unbiased estimate of this combined risk is

$$afic^u(M) = \int \{\widehat{\sigma}_M(v)^2/n\} dW(v) + \int \{\widehat{b}_M(v)^2 - \widehat{\kappa}_M(v)^2/n\} dW(v).$$

This is the same as a direct weighted sum or integral of the individual $FIC^u(M, v)$ scores. The adjusted version, however, where a potentially negative value of the estimated integrated squared bias is being truncated to zero, is not identical to the integral of the $FIC(M, v)$ scores. It is rather equal to

$$afic(M) = \int \{\widehat{\sigma}_M(v)^2/n\} dW(v) + \max\left[\int \{\widehat{b}_M(v)^2 - \widehat{\kappa}_M(v)^2/n\} dW(v), 0\right].$$

As with FIC, there are two related, but not identical, approximation schemes, the fixed wide model setup and the local asymptotics, of respectively section 3.1 and **Appendix A1**, leading now to somewhat different AFIC formulae. For details and applications (see Claeskens and Hjort, 2008a,b, Ch. 6).

There is a connection between Akaike's information criterion AIC and AFIC with certain model dependent weights (see Claeskens and Hjort, 2008a, Sec. 6.2). Broadly speaking, the AIC turns out to be large-sample equivalent to cases with AFIC where "all things are equally important."

3. FIC WITHIN A FIXED WIDE MODEL FRAMEWORK

The FIC as used in the bird species example is the version as derived in Claeskens and Hjort (2003), see also Claeskens and Hjort (2008b, Ch. 6). For the estimation of bias and variance a local asymptotic framework is used in which the parameters of the true density of the data are assumed to be of the form $\gamma = \gamma_0 + \delta/\sqrt{n}$, with n the sample size, see **Appendix A1** for more explanation. This assumptions means in practice that we believe that all models are relative close to each other and to the truth. Moreover, all models are submodels of a wide model. Since the derivation of the FIC formulae is contained in the references above, we only place a summary in the **Appendix**.

In this section we present the "fixed wide model" framework, which is particularly useful if the set of candidate models are seen as not being in a reasonable vicinity of each other. This second framework allows candidate models of a different sort from the wide model; in particular, a candidate model does not have to be a clear submodel of the wide model. Keep in mind that the two different FIC frameworks have the same aims and motivation; the difference between them lies in the different mathematical tools for estimating the relevant mse quantities, which lead to different formulae. In the discussion section 5 we come back to some differences between the two frameworks. Here we start in section 3.1 by presenting the fixed wide model FIC in a general regression setup. Then in the two following subsections we deal with two specific model classes of general interest, generalized linear models and linear mixed models, in more detail.

3.1. General Regression Models

In this subsection we use the familiar (x_i, y_i) notation for the regression data, with x_i the covariate vector in question. The FIC machinery we develop here starts from the existence of a fixed wide model. The development represents an extension of earlier work of Jullum and Hjort (2017, 2019) for i.i.d. data and survival analysis, Ko et al. (2019) for copulae models, Cunen et al. (2019) for power-law distributions (with applications to war and conflict data) and Cunen et al. (in review)^{1,2} for linear mixed effects models (with application to whale ecology).

¹Cunen, C., Walløe, L., and Hjort, N. L. (2019). Focused model selection for linear mixed models, with an application to whale ecology. *Ann. Appl. Stat.*

²Cunen, C., Walløe, L., Konishi, K., and Hjort, N. L. (2019). Decline in energy storage for the Antarctic minke whale (*Balaenoptera bonaerensis*) in the Southern Ocean during the 1990s.

Since we wish to estimate the mse of the focus estimator in different models, we first consider the asymptotic distribution of the parameter estimator in the wide model and next in the other models of interest. The distributions are used to form the mse's of the focus estimators and finally we construct the fic as an estimated mse and select the model with the smallest fic value.

Suppose a wide model density is agreed upon, of the form $f(y_i | x_i, \theta)$, for a certain parameter vector θ , of length p . We consider this to be the true model. This θ would typically encompass both regression coefficients and parameters related to the spread and shape of error distributions. Define $u(y_i | x_i, \theta) = \partial \log f(y_i | x_i, \theta) / \partial \theta$ the score function, and $J_n = n^{-1} \sum_{i=1}^n \text{Var}_{\text{wide}} u(Y_i | x_i, \theta_{\text{true}})$ the normalized Fisher information matrix at the true parameter. Under mild regularity conditions we have the following well-known result for the maximum likelihood estimator $\hat{\theta}_{\text{wide}}$,

$$\sqrt{n}(\hat{\theta}_{\text{wide}} - \theta_{\text{true}}) \approx_d N_p(0, J_n^{-1}). \tag{5}$$

The notation indicates approximate multinormality to the first order as the sample size grows, and can also be supplemented with a clear limit distribution statement, in that case involving a limit covariance matrix J for J_n . Consider now a candidate model M , different from the wide one, perhaps also in structure and form. With notation $f_M(y_i | x_i, \theta_M)$ for its density, and $u_M(y | x_i, \theta_M)$ for its score function, we have a maximum likelihood estimator $\hat{\theta}_M$, of length p_M , maximizing the log-likelihood function $\ell_{n,M}(\theta_M) = \sum_{i=1}^n \log f_M(y_i | x_i, \theta_M)$. If the wide model is considered to be the truth, the estimator in model M does not necessarily aim at the true parameter, but at the least false parameter $\theta_{0,M,n}$, which is the minimizer of the Kullback–Leibler distance from the data-generating mechanism to the model; see details in **Appendix A3**. The estimator in the candidate model has a limiting multinormal distribution, with a sandwich type variance matrix,

$$\sqrt{n}(\hat{\theta}_M - \theta_{0,M,n}) \approx_d N_{p_M}(0, J_{M,n}^{-1} K_{M,n} J_{M,n}^{-1}), \tag{6}$$

where

$$J_{M,n} = -n^{-1} \sum_{i=1}^n E_{\text{wide}} \frac{\partial^2 \log f(Y_i | x_i, \theta_{0,M,n})}{\partial \theta_M \partial \theta_M^t} \quad \text{and}$$

$$K_{M,n} = n^{-1} \sum_{i=1}^n \text{Var}_{\text{wide}} u_M(Y_i | x_i, \theta_{0,M,n}).$$

The variance matrices here are defined with respect to the wide model, at position θ_{true} .

From approximations (5–6) the delta method may be called upon to read off relevant expressions for the approximate distributions of the focus parameter estimators $\hat{\mu}_{\text{wide}} = \mu(\theta)$ and $\hat{\mu}_M = \mu_M(\theta_M)$, where the latter is aiming for the least false parameter value $\mu_{0,M,n} = \mu_M(\theta_{0,M,n})$ associated with model M . Crucially, we also need a multinormal approximation to the *joint* distribution of $(\hat{\mu}_{\text{wide}}, \hat{\mu}_M)$, in order to assess the distribution of the bias estimator $b_M = \hat{\mu}_M - \hat{\mu}_{\text{wide}}$; without that part we can't

build an appropriate estimator for b_M^2 . In **Appendix A3**, we go through such arguments, and reach

$$\begin{pmatrix} \sqrt{n}(\hat{\mu}_{\text{wide}} - \mu_{\text{true}}) \\ \sqrt{n}(\hat{\mu}_M - \mu_{0,M,n}) \end{pmatrix} \approx_d N_2(0, \Sigma_{M,n}). \tag{7}$$

Here the 2×2 matrix $\Sigma_{M,n}$ has diagonal terms $c^t J_n^{-1} c$ and $c_{M,n}^t J_{M,n}^{-1} K_{M,n} J_{M,n}^{-1} c_{M,n}$, with gradient vectors

$$c = \partial \mu(\theta_{\text{true}}) / \partial \theta \quad \text{and} \quad c_{M,n} = \partial \mu(\theta_{0,M,n}) / \partial \theta_M$$

of lengths p and p_M . The off-diagonal term of $\Sigma_{M,n}$ takes the form $c^t J_n^{-1} C_{M,n} J_{M,n}^{-1} c_{M,n}$, with a formula for the required covariance related term $C_{M,n}$ in the **Appendix**.

From (7) we can read off mse approximations,

$$\begin{aligned} \text{mse}_{\text{wide}} &\doteq c^t J_n^{-1} c / n + 0^2 \quad \text{and} \\ \text{mse}_M &\doteq c_{M,n}^t J_{M,n}^{-1} K_{M,n} J_{M,n}^{-1} c_{M,n} + b_M^2, \end{aligned}$$

with bias $b_M = \mu_{0,M,n} - \mu_{\text{true}}$. For the latter we use the estimator $\hat{b}_M = \hat{\mu}_M - \hat{\mu}_{\text{wide}}$, where the result above also leads to a clear approximation for the distribution of $\sqrt{n}(b_M - \hat{b}_M)$. This leads to FIC formulae, unbiased and adjusted, as

$$\begin{aligned} \text{FIC}_{\text{wide}}^u &= \hat{c}^t \hat{J}_n^{-1} \hat{c} / n + 0^2 \quad \text{and} \\ \text{FIC}_M^u &= \hat{c}_M^t \hat{J}_M^{-1} \hat{K}_M \hat{J}_M^{-1} \hat{c}_M / n + \hat{b}_M^2 - \hat{\kappa}_M^2 / n, \\ \text{FIC}_{\text{wide}} &= \hat{c}^t \hat{J}_n^{-1} \hat{c} / n + 0^2 \quad \text{and} \\ \text{FIC}_M &= \hat{c}_M^t \hat{J}_M^{-1} \hat{K}_M \hat{J}_M^{-1} \hat{c}_M / n + \max(\hat{b}_M^2 - \hat{\kappa}_M^2 / n, 0). \end{aligned} \tag{8}$$

Here \hat{c} and \hat{c}_M emerge by computing gradients of $\mu(\theta)$ and $\mu_M(\theta_M)$ at their respective maximum likelihood positions, and \hat{J}_n, \hat{J}_M are computed as normalized observed Fisher information matrices, for the wide and for the candidate model in question; specifically, \hat{J}_M is $1/n$ times minus the Hessian matrix from the log-likelihood, $-\partial^2 \ell_{n,M}(\hat{\theta}_M) / (\partial \theta_M \partial \theta_M^t)$. Also, the $p_M \times p_M$ matrix \hat{K}_M is $n^{-1} \sum_{i=1}^n \hat{u}_{M,i} \hat{u}_{M,i}^t$, with $\hat{u}_{M,i} = u_M(y_i | x_i, \hat{\theta}_M)$. Finally, the $\hat{\kappa}_M^2 / n$ estimates involves also the $p \times p_M$ matrix \hat{C}_M , which is $n^{-1} \sum_{i=1}^n \hat{u}_{\text{wide},i} \hat{u}_{M,i}^t$. Model selection proceeds by computing FIC_M , the estimated mse of the focus estimator $\hat{\mu}_M$, for all models M of interest, and then selecting that model for which this score is the lowest.

3.2. FIC for Generalized Linear Models, With a Fixed Wide Model

We illustrate this FIC machinery for one popular class of generalized linear models, namely the Poisson regression models. Generalizations to other generalized linear models are relatively immediate. Suppose therefore that we have count data y_i along with a covariate vector x_i of length p . For the fixed wide model we take the Poisson regression model with $y_i \sim \text{Pois}(\xi_i)$, with $\xi_i = \exp(x_i^t \beta)$ containing all covariate information; in particular, there is also a true parameter β_{true} there. Consider then an alternative candidate model M which instead takes the means to be $\xi_{M,i} = \exp(x_{M,i}^t \beta_M)$, with $x_{M,i}$ of length p_M , perhaps a subset of the full x_i , or perhaps with some entirely other pieces

of covariate information. Here the log-densities take the form $-\xi_i + y_i \log \xi_i - \log(y_i!)$, which means

$$\begin{aligned} \log f &= -\exp(x_i^t \beta) + y_i x_i^t \beta - \log(y_i!) \quad \text{and} \\ \log f_M &= -\exp(x_{M,i}^t \beta_M) + y_i x_{M,i}^t \beta_M - \log(y_i!), \end{aligned}$$

for the wide model and the candidate model, along with score functions

$$\begin{aligned} u(y_i | x_i, \beta) &= \{y_i - \exp(x_i^t \beta)\} x_i \quad \text{and} \\ u_M(y_i | x_{M,i}, \beta_M) &= \{y_i - \exp(x_{M,i}^t \beta_M)\} x_{M,i}. \end{aligned}$$

From this we deduce

$$\begin{aligned} J_n &= n^{-1} \sum_{i=1}^n \exp(x_i^t \beta_{\text{true}}) x_i x_i^t, \\ J_{M,n} &= n^{-1} \sum_{i=1}^n \exp(x_{M,i}^t \beta_{0,M,n}) x_{M,i} x_{M,i}^t, \\ K_{M,n} &= n^{-1} \sum_{i=1}^n \exp(x_i^t \beta_{\text{true}}) x_{M,i} x_{M,i}^t, \end{aligned}$$

along with the $p \times p_M$ covariance matrix $C_{M,n}$, defined as

$$\begin{aligned} n^{-1} \sum_{i=1}^n E_{\text{wide}} \{Y_i - \exp(x_i^t \beta_{\text{true}})\} x_i \{Y_i - \exp(x_{M,i}^t \beta_{0,M,n})\} x_{M,i}^t \\ = n^{-1} \sum_{i=1}^n \exp(x_i^t \beta_{\text{true}}) x_i x_{M,i}^t. \end{aligned}$$

Consistent estimates of these population matrices are obtained by inserting $\hat{\beta}_{\text{wide}}$ for β_{true} and $\hat{\beta}_M$ for $\beta_{0,M,n}$.

Notably, as long as there is a well-defined wide Poisson regression model, as assumed here, the framework is sufficiently flexible and broad to encompass also non-Poisson candidate models. Using the FIC apparatus involves working with log-likelihood functions and score functions for these alternative models, leading to different but workable expressions for the matrices $J_{M,n}$, $K_{M,n}$, $C_{M,n}$ above. The stretched Poisson models used in Schweder and Hjort (2016, Exercise 8.18) are a case in point; these allow both over- and underdispersion.

3.3. FIC for Linear Mixed Effects Models

Models with random effects, often called mixed effect models, are widely used in ecological applications. In Cunen et al. (in review)¹ FIC formulae have been developed for the class of linear mixed effect models (often abbreviated LME models). Here we will give a brief description of that approach, which also serves as a special case of the general FIC approach for a fixed wide model framework, see (8). Generalizations to classes of non-linear mixed effect models, and also to heteroscedastic situations where variance parameters depend on covariates, can be foreseen, following similar chains of arguments but involving more elaborations.

Suppose we have n observations of y_i , a vector of length m_i . The m_i datapoints within each y_i vector are assumed to be

dependent, and will often correspond to data collected in the same space or time. Here we will refer to these data as belonging to the same *group*. Each y_i vector is associated with a regressor matrix X_i of dimension $m_i \times p$ for the fixed effects, and a design matrix Z_i of dimension $m_i \times k$ for the random effects. The linear mixed effects model takes the form

$$y_i = X_i \beta + Z_i b_i + \varepsilon_i \quad \text{for } i = 1, \dots, n,$$

with the $b_i \sim N_k(0, D)$ independent of the errors $\varepsilon_i \sim N_{m_i}(0, \sigma^2 I_{m_i})$. The model may also be represented as

$$Y_i \sim N_{m_i}(X_i \beta, \sigma^2 (I_{m_i} + Z_i D Z_i^t)), \tag{9}$$

and its parameters are $\theta = (\beta, \sigma, D)$. Note that the ordinary linear regression model is a special case, corresponding to $D = 0$. The log-likelihood contribution for this group of the data may be written

$$\begin{aligned} \ell_i(\theta) &= -m_i \log \sigma - \frac{1}{2} \log |I_{m_i} + Z_i D Z_i^t| - \frac{1}{2} (1/\sigma^2) (y_i - X_i \beta)^t \\ &\quad \times (I_{m_i} + Z_i D Z_i^t)^{-1} (y_i - X_i \beta). \end{aligned}$$

The combined log-likelihood $\sum_{i=1}^n \ell_i(\theta)$ leads to maximum likelihood estimators and hence also to $\hat{\mu}_{\text{wide}} = \mu(\hat{\beta}_{\text{wide}}, \hat{\sigma}_{\text{wide}}, \hat{D}_{\text{wide}})$ for any focus parameter $\mu = \mu(\beta, \sigma, D)$ of interest.

In applied situations we will spend efforts and call on biological knowledge to construct a well-motivated wide model, of the form (9). This wide model will typically be based on our knowledge of the system under study and, crucially, on how the data were collected. Quite often the resulting model could become *big*, in the sense that it includes a large number p of fixed effects and also a large number k of random effects. Assume, as we do throughout this paper, that our primary interest lies in the precise estimation of some focus parameter μ , which could be a function of the fixed effect coefficients β , and/or the variance components (σ, D) . For such a $\mu = \mu(\beta, \sigma, D)$, can we find another model which offers more precise estimates of μ than $\hat{\mu}_{\text{wide}} = \mu(\hat{\beta}_{\text{wide}}, \hat{\sigma}_{\text{wide}}, \hat{D}_{\text{wide}})$ implied by the wide model?

FIC answers the question above; we can search among a set of candidate models for one giving more precise estimates of μ . In the simplest setting, the candidate model is defined with respect to the same n groups as in the wide model in (9), and we write

$$y_i \sim N_{m_i}(X_{M,i} \beta_M, \sigma_M^2 (I + Z_{M,i} D_M Z_{M,i}^t)).$$

This model has design matrices, $X_{M,i}$ and $Z_{M,i}$, potentially different from those of the wide model, and hence also a different set of parameters, say $\theta_M = (\beta_M, \sigma_M, D_M)$. Often, but not necessarily, the candidate model will involve subsets of the covariates (i.e., columns) included in X_i and Z_i , respectively. Let the covariate matrix $X_{M,i}$ have dimension $m_i \times p_M$, and $Z_{M,i}$ being $m_i \times k_M$. The focus parameter must then be represented properly inside the candidate model, as $\mu_M = \mu_M(\beta_M, \sigma_M, D_M)$, leading to the estimate $\hat{\mu}_M = \mu_M(\hat{\beta}_M, \hat{\sigma}_M, \hat{D}_M)$.

In order to work out FIC formulae, we first need to study the joint large-sample behavior of the estimator from the wide model

$\widehat{\mu}_{\text{wide}}$ and the estimator from the candidate model $\widehat{\mu}_M$. This is as with Equation (7) in section 3.1, but the current framework is more complicated and needs further efforts. Such work is carried out in Cunen et al. (in review)¹, and lead to

$$\left(\frac{\sqrt{n}(\widehat{\mu}_{\text{wide}} - \mu_{\text{true}})}{\sqrt{n}(\widehat{\mu}_M - \mu_{0,M,n})} \right) \approx_d N_2(0, \Sigma_{M,n}),$$

with all quantities defined analogously to what is presented in section 3.1. These include matrices $J_n, J_{M,n}, K_{M,n}, C_{M,n}$ and gradient vectors c and $c_{M,n}$, defined similarly to those in section 3.1, but here involving more complicated details than for the plainer regression models worked with there.

This work then yields the same type of FIC formulae as for Equation (8), but with other recipes and formulae for the required estimators for the quantities mentioned. Regarding estimators for the matrices involved, we have three general possibilities: (i) working out explicit formulae and plug in the necessary parameter estimates; (ii) computing the matrices numerically, involving certain numerical integration details; (iii) via bootstrapping from the estimated wide model. In Cunen et al. (in review)¹ the first option is pursued, involving lengthy derivations of log-density derivatives and their means, variances, covariances, computed under the wide model. The resulting formulae are too long for this review, but are fast to compute. Options (ii) and (iii) have yet to be fully investigated, but will likely be fruitful when extending this FIC approach to the wider class of generalized linear mixed models (the so-called GLMMs).

The approach described here will be illustrated in section 4, but we first offer some comments of a more general nature. Readers familiar with linear mixed effects models will be aware that there are two different estimation schemes for models of this class, full maximum likelihood and so-called REML estimators, for restricted or residual maximum likelihood. The REML method takes the estimation of the fixed effects of the model into account when producing estimators of the variance parameters. For the computation of FIC scores the user might employ either maximum likelihood or residual maximum likelihood estimates, since these are large-sample equivalent; see for instance Demidenko (2013, Ch. 3). As with the general FIC formulae (8) there are two versions, the approximately unbiased estimates of risks and the adjusted ones. In Cunen et al. (in review)¹ it is argued that the unbiased version

$$FIC_M^u = \widehat{c}_M^t \widehat{J}_M^{-1} \widehat{K}_M \widehat{J}_M^{-1} \widehat{c}_M / n + \widehat{b}_M^2 - \widehat{\kappa}_M^2 / n \quad (10)$$

tends to work best for linear mixed effects models. The benefit of this version is that good candidate models with small biases earn more, compared to the wide model. Investigations show that the FIC formulae of (10) work well, in the sense that they accurately estimate the risk associated with the use of the different candidate models. The FIC formulae are based on large-sample arguments, which for the case of the linear mixed effects models involves approximations to normality when the number n of groups increases to infinity. These normal approximations work well as long as the full sample size $\sum_{i=1}^n m_i$ grows, particularly for functions of the linear mean parameters. More care is

sometimes required when it comes to applications involving non-linear functions of both mean and variance parameters, as with estimates of probabilities $\mu = P(Y \geq y_0 | x_0, z_0)$.

4. APPLICATION: THE SLIMMING OF MINKE WHALES

Our second application story concerns the potential change in body condition of Antarctic minke whales over a period of 18 years. For a more thorough investigation consult Cunen et al. (in review)². Questions treated there have been discussed in the Scientific committee of the International Whaling Commission (IWC) for a number of years, and a full consensus has not been reached. In the context of this review, therefore, the analysis below should be taken as an illustration, and not necessarily the last word on the topic of the decline in energy storage or body condition for the minke whales.

Using data from the Japanese Whale Research Program under Special Permit in the Antarctic (the so-called JARPA-1) we have studied the evolution of fat weight in Antarctic minke whales caught in 18 consecutive years, from 1988 and 2005. The main biological interest lies in whether or not the whales experienced a decline in body condition during the study period, and the dissected fat weight (in tons or kg) is taken to be a proxy for this body condition. Thus, there is a clear focus parameter in this application: the yearly decline in fat weight (which we will parametrize in a suitable fashion in the following).

The whales caught in each year are unevenly sampled with respect to a number of covariates, for instance sex, body length, age, and longitudinal region in the Antarctic ocean. Since all these covariates may influence body condition we need to include them in a model aiming at estimating the potential yearly decline in the response. Based on lengthy and detailed discussions in the Scientific Committee of the IWC, we have chosen a wide model within the class of linear mixed effect models, see section 3.3. In Cunen et al. (in review)² we have used considerable efforts to motivate the choice of covariates, interactions, and random effect terms in the wide model, but these arguments are outside the scope of the present article. In R-package-type notation, the wide model can be given as

$$\begin{aligned} \text{fatweight} \sim & \text{year} + \text{year}^2 + \text{bodylength} + \text{sex} \\ & + \text{diatom} + \text{date} + \text{date}^2 + \text{age} \\ & + \text{sex} * \text{diatom} + \text{diatom} * \text{date} \\ & + \text{diatom} * \text{date}^2 + \text{bodylength} * \text{sex} \\ & + \text{bodylength} * \text{date} \\ & + \text{bodylength} * \text{date}^2 + \text{sex} * \text{date} \\ & + \text{sex} * \text{date}^2 \\ & + \text{bodylength} * \text{sex} * \text{date} \\ & + \text{bodylength} * \text{sex} * \text{date}^2 \\ & + \text{age} * \text{sex} \\ & + \text{age} * \text{date} + \text{age} * \text{date}^2 \\ & + \text{age} * \text{sex} * \text{date} + \text{age} * \text{sex} * \text{date}^2 \\ & + \text{year} * \text{sex} + \text{year}^2 * \text{sex} + \text{region} \end{aligned}$$

+ year * region + year² * region
 + sex * region + diatom * region
 + region * date
 + region * date² + (1 + date
 + date² | year).

The region covariate reflects three different geographical regions, associated with three regression coefficients summing to zero.

The model defined above has $p = 40$ fixed effect coefficients. The notation $(1 + date + date^2 | year)$ specifies the random effect structure; the groups are defined by a categorical version of the year variable (so $n = 18$), and the Z_i matrix has $k = 3$ columns (a column of ones for the intercept, date, and date squared). According to prior biological knowledge, date is assumed to be one of the most important effects governing the fat weight. The variable refers to the day of the season when each whale was caught, and since the whales are in the Antarctic to gain weight the coefficient related to date is expected to be large and positive. Also, the effect of date is expected to be different from year to year, possibly due to fluctuations in krill production. Hence, a random effect on date is included. We thus have a total of $40 + 1 + 6 = 47$ parameters to estimate. The total number of observations, i.e., $\sum_{i=1}^n n_i$, was 683.

As mentioned above the main interest, for discussions at several IWC meetings, has been the yearly decline in the fatweight outcome variable. Since we have a quadratic year term in our wide model, with that part taking the form $\beta_{year}x + \beta_{year2}x^2$ for year x , a natural definition of the yearly decline is $\mu = \beta_{year} + 2\beta_{year2}x_0$, with x_0 the mean year in the dataset. The focus parameter corresponds to the derivative of the mean response, with respect to year, and evaluated in this mean year time point. For candidate models with only a linear effect of year the parameter simplifies to β_{year} only. Furthermore, for those submodels where there is no year effect included, we have $\beta_{year} = 0$, a parameter value which then is estimated with zero variance but with potentially big bias. For this example, we have limited ourselves to investigating five candidate models only, in addition to using the wide model itself; see **Table 2**.

We do not actually expect the mean level of decline in energy storage to be either exactly linear or exactly quadratic over 18 years, but take this level of approximation to be adequate for the purpose, since the decline over time curve is not far from zero; also, our focus parameter is identical to the overall slope, the mean curve evaluated at the end point minus its value at the start point, divided by the length of time.

All the candidate models have a smaller number of fixed effects than the wide model. Note that the first candidate model M_1 has a more complex random effect structure than the wide model itself (with $k = 6$ giving a total of 21 random effect parameters). This choice also demonstrates that there is nothing in the formulae hindering us from having candidate models with more random effects (or also more fixed effects) than the wide model. When it comes to interpreting the results, it is usually more natural to choose the wide model to be the largest possible plausible model, however. The models M_2 and M_3 are very simple (with few fixed

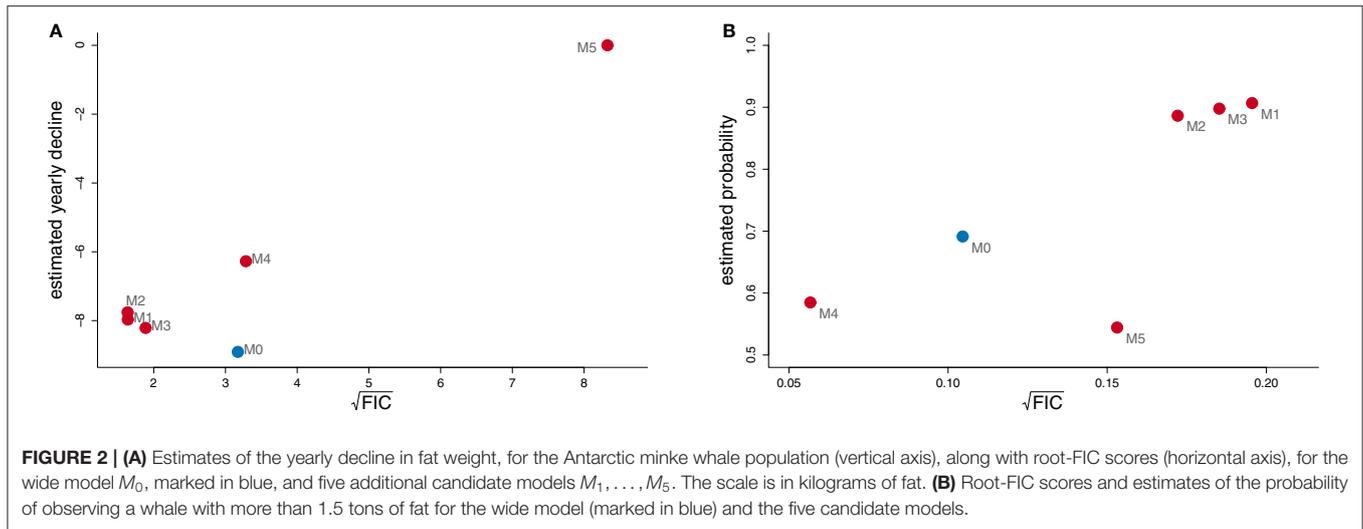
TABLE 2 | Brief descriptions of the wide model and the five additional candidate models, with the number of fixed effects, the number of random effects, and the total number of parameters to be estimated, for each model.

	Description	p	k	d
M_0	Wide model	40	3	47
M_1	Less interactions, quadratic year effect	9	6	31
M_2	Very simple, linear year effect	5	2	9
M_3	Very simple, linear year effect	5	1	7
M_4	Only linear year effect	2	1	4
M_5	Like the wide, but without year effect	32	3	39

effects), and differ only in their random effects. Model M_4 includes only the linear year effect in addition to a single random effect in the intercept. The last model, M_5 , is the model without any year effect, so $\mu_{M_5} = 0$. With the present focus parameter, the FIC score of such a model will have zero variance and a bias which only depends on the estimated focus parameter in the wide model, and its estimated variance, so $FIC_{M_5}^u = (0 - \hat{\mu}_{wide})^2 - \hat{\kappa}_{wide}^2/n$, for the relevant κ_{wide}^2/n approximation to the variance of $\hat{\mu}_{wide}$. Thus, further specification of M_5 is unnecessary; it includes all possible LME models without any year effect. As the candidate models worked with are not close enough to each other to warrant the use of the local neighborhoods framework, we use the “fixed wide model” approach.

After carefully constructing our wide model, and checked that it passes various diagnostic tests, we can proceed to model selection with the FIC. The results are given in the form of a FIC-plot in **Figure 2**. We see that M_2 gets the lowest FIC score, with $\hat{\mu} = -7.76$. The models M_1 and M_3 are close to the winning one, both in terms of their FIC scores and their estimates of the focus parameter. Model M_5 , without the year effect had a considerably larger FIC score than any of the other models (which can be seen as an implicit test for the the null hypothesis of there being no year effect). From the plot we can conclude that our best estimate of the focus parameter is around -8 kg, or 80 kg loss of fat over a decade. Furthermore, since the root-FIC values are about 1.50, confidence intervals associated with these best point estimates will clearly fall to the left of zero. A natural interpretation of the FIC plot is therefore that the body condition decline, for the Antarctic minke whales, has been *negative and significant* over the study period.

To demonstrate the versatility of our approach, we have investigated the same six models with respect to another focus parameter, the probability of observing a whale with more than a certain amount of fat, say 1.5 tons (1,500 kg), given some covariate values: $\mu_2 = P(Y \geq 1.5 | x_0, z_0)$. Here we chose to look at a 20 year old male whale, caught in 1991 in the eastern region, of approximately mean length (8 m), and which is caught toward the end of the season. Over the full dataset, the average fat weight of a whale is close to 1.5 tons. The FIC scores and estimates are given in **Figure 2**. We observe that the models give widely different estimates, ranging from around 0.50 to 0.90, and that the ranking of the models is very different from the ranking when the focus was the yearly decline in fat weight. The smallest model M_4 is considered the best for estimating the probability of



observing a “medium fat” whale. Here, we see the typical bias-variance trade-off at work: using M_4 clearly gives an estimate with some bias compared to the wide model (estimate of 0.60 instead of around 0.70), but the bias is compensated for by a strong decrease in variance.

5. DISCUSSION

Our article has motivated, exhibited, developed, and extended the machinery of Focused Information Criteria for model selection and model ranking, with a few illustrations for ecological data. Here we offer some general remarks.

1. *The role of the wide model.* The FIC idea is to examine how different candidate models work regarding what they actually deliver, in terms of point estimates for the most crucial parameters of interest. This examination involves approximations to and estimates of the risks, which for the usual squared error loss function means mean squared error. Quantifying the implied variances and biases relies on the notion of a clearly defined (though unknown) data generating mechanism. This is one of the roles of our *wide model*. In the local asymptotics framework of **Appendix A1** this is the full model $f(y_i | x_i, \theta, \gamma)$ of (12), with $p + q$ parameters; in the alternative framework of section 3.1 it is what we term the fixed wide model. Such a wide model needs to be well argued, as being sufficiently rich to encompass the anticipated salient features of the phenomena studied. Since quantification and consequent estimation of variances and biases rest on the wide model being adequate it ought also to be given a goodness-of-fit verification, involving diagnostic checks etc.

One might inquire how sensitive the FIC scores are to the choice of the wide model. In connection with the application described in section 4 we have conducted some sensitivity checks and found that moderate changes to the wide model had little effect on the ranking of the different candidate models. Also, for the wide models we have investigated, the estimate of the focus parameter in the selected models was reasonably stable.

More radical changes to the wide model should be expected to have greater effect, but we have not fully investigated this issue. Fully guarding against all misspecification of the wide model is unattainable, but extending our approach to even wider and more flexible wide models may lead to some improvements.

2. *When should you use FIC?* Practitioners may be interested in model selection for different, overlapping reasons. On one hand the goal might be to select the candidate model which in a relevant sense is the closest to the true data generating mechanism. Criteria based on model fit and some penalization for complexity aim at this goal, for instance the well-known AIC and BIC (see Claeskens and Hjort, 2008b) for a general discussion. On the other hand, practitioners often seek a small model offering precise estimates of the quantities they are interested in. It is important to keep in mind that FIC specifically aims at the second goal, and is not necessarily suitable for the first goal. FIC offers a principled way to *simplify* a large, realistic model which the user assumes to hold (i.e., to be realistically and adequately close to the complicated truth). The goal of the simplification is to obtain more precise estimates of quantities of interest, say $\hat{\mu}$ for an underlying focus parameter μ . This also includes producing predictions for not yet seen outcomes of random variables, like the abundance of a certain species over the coming twenty years. Here simplification must be understood in a wide sense, as the candidate models do not necessarily need to be nested within the wide model, as we have seen. The two different motivations for model selection alluded to above partly relate to the two goals for statistical modeling: to explain or to predict, i.e., the “two cultures of statistics” (see Breiman, 2001; Shmueli, 2010). For yet further perspectives on model selection with focused views, coupled with model structure adequacy analysis (see Taper et al., 2008).

Once a practitioner has decided to use FIC, she then has to make a choice between the two FIC frameworks we have discussed, using local asymptotics or a fixed wide model. As a tentative guiding rule we advocate turning to the “fixed wide model” setup if the set of candidate models are seen as not being in a reasonable vicinity of each other. Also, we have seen that this

framework allows candidate models of a different sort from the wide model; in particular, a candidate model does not have to be a clear submodel of the wide model. As stated before, the two frameworks aim at the same quantities, and the choice may thus also be guided by convenience. Note also that in many situations the two frameworks may give similar results. For the special case of linear regression models with focus parameters being linear functions of the coefficients, the formulae turn out to be identical. Also, for the classical generalized linear models, including logistic and Poisson regressions, the formulae yield highly correlated scores, as long as the focus parameters under study are functions of such linear combinations $x_0^t \beta + z_0^t \gamma$. For more complicated focus parameters, like probabilities for crossing thresholds, the answers are not necessarily close, and will depend on both the sample size and the degree to which the candidate models are not close.

3. *Model averaging.* Model averaging is sometimes used as an alternative to model selection to avoid the perhaps brutal throwing away of all but one model. With model averaging one computes the estimate of the focus quantity in all of the models separately and then forms a weighted average which is used as the final “model averaged” estimate of the focus. See for example the overview paper about model averaging in ecology by Dormann et al. (2018). Averaging estimates has as the advantage that all models are used. The flexibility of choosing the weights allows to give a larger weight to the estimate of a model that one prefers most. Weights could be set in a deterministic way, such as giving equal weights to all estimates, or could be data-driven. It makes sense to use values of information criteria to set the weights. Especially AIC has been popular (see Burnham and Anderson, 2002) for examples of the use of “Akaike weights.” Also FIC could be used to form weights that are proportional to $\exp(-\lambda \text{FIC}_M / \text{FIC}_{\text{wide}})$ for a user-chosen value of λ . One could also try to set the weights such that the mean squared error of the weighted estimator is as small as possible (Liang et al., 2011). Such theoretically optimal weights need to be estimated for practical use, which induces again estimation variability, and might lead to a more variable weighted estimator as when simple equal weights would have been used (Claeskens et al., 2016).

Model averaging with data-driven weights has consequences for inference similar to the post-selection inference (see below). Indeed, model selection may be seen as a form of model averaging, with all but one of the weights equal to zero and the remaining weight equal to one. Correct frequentist inference for model averaged estimators needs to take the correlations between the separate estimators into account, as well as the randomness of the weights in case of data-driven weights.

4. *Post-selection issues.* Model selection by the use of an information criterion (such as FIC, or AIC) comes with several advantages as compared to contrasting models two by two via hypothesis testing. With model selection there is no need to single out one model that would be placed in a null hypothesis. All models are treated equally. Multiple testing issues do not occur because no testing takes place. The set of models that is searched over can be large. The ease of calculating such information criteria makes it fast and allows to include many models in

the search. However, there is a price to pay when one puts the next step to perform inference using the selected model. Simply ignoring that a model is arrived at via a selection procedure results in p-values that are too small and confidence intervals that are too narrow.

With a replicated study resulting in a dataset similar to but independent of the current one, it might happen that a different model gets selected, all the rest left unchanged. This illustrates that variability is involved in the process of model selection. One way to address such variability is via model averaging (see e.g., Hjort and Claeskens, 2003, Claeskens and Hjort, 2008b, Ch. 7, Efron, 2014). Berk et al. (2013) develop an approach for the construction of confidence intervals for parameters in a linear regression model that uses a selected model. Their approach is conservative, in the sense that the intervals tend to be wide and sometimes have a coverage that is quite a bit larger than the nominal value. Other approaches to take the uncertainty induced by the selection procedure into account is via selective inference leading to so-called “valid” inference. See, for example, Tibshirani et al. (2016, 2018). By using information about the specifics of the selection method such inference methods result in narrower confidence intervals as compared to the Berk et al. (2013) method. The effect of increasing the number of models results in getting larger confidence intervals (see Charkhi and Claeskens, 2018). Valid inference after selection is currently investigated for several model selection methods. It is to be expected that more results will become available in the future that guarantee that working with a selected model happens in a honest way that takes all variability into account.

It is well-known that estimators computed under a given model become approximately normal, under mild regularity conditions. It is however clear from the brief discussion above that post-selection and more general model-average estimators have more complicated distributions, as they often are non-linear mixtures of approximately normal distributions, with different biases, variances, and correlations. Clear descriptions of large-sample behavior, for even complex model-selection and model-average schemes, can be given inside the local asymptotics $O(1/\sqrt{n})$ framework of **Appendix A1**, as shown in Hjort and Claeskens (2003), Claeskens and Hjort (2008b, Ch. 7), with further generalizations in Hjort (2014). Inside the general framework of (12), with estimators $\hat{\mu}_M$ as in (13), consider the combined or post-selection estimator

$$\hat{\mu}^* = \sum_M \hat{w}(M) \hat{\mu}_M,$$

with data-driven weights $\hat{w}(M)$ summing to one. If these are weights take the form $w(M | D_n)$, with $D_n = \sqrt{n}(\hat{\gamma}_{\text{wide}} - \gamma_0)$ as in (15), there is a very clear limit distribution,

$$\sqrt{n}(\hat{\mu}^* - \mu_{\text{true}}) \rightarrow_d \Lambda_0 + \omega^t \{ \delta - \hat{\delta}(D) \}, \quad \text{where} \\ \hat{\delta}(D) = \sum_M w(M | D) G_M D. \quad (11)$$

This extends the master theorem result (17), to allow even for very complicated post-selection and model averaging schemes.

The $q \times q$ matrices G_M in this orthogonal decomposition are as in (16). The result remains true also for schemes based on weights involving AIC or FIC weights, as the appropriate weights can be shown to be close enough to the relevant $w(M|D_n)$. These limiting distributions can be simulated, at any position in the δ domain. Yet further efforts are required to turn such into valid post-selection or post-averaging confidence intervals, however see Claeskens and Hjort (2008b, Ch. 7) for one particular general (conservative) recipe, and for further discussion of these issues.

5. *Performance.* It is beyond the scope of this article to go into the relevant aspects of statistical performance of the FIC methods. One may indeed study both the accuracy of the final post-selection or post-averaged estimator, say for the $\hat{\mu}^*$ above, and the probabilities for selecting the best models. Such questions are to some extent discussed in Hjort and Claeskens (2003) and Claeskens and Hjort (2008b, Ch. 7); broadly speaking, the FIC outperforms the AIC in large parts of the parameter space, but not uniformly. There are also several advantages with FIC, when compared with the BIC, regarding precision of the finally evaluated estimators. Notably, all of these questions can be studied accurately in the limit experiment alluded to above, where all limit distributions can be given in terms of the orthogonal decomposition $\Lambda_0 + \omega^t\{\delta - \hat{\delta}(D)\}$ of (11).

6. *FIC for high-dimensional data.* When models contain a large number of parameters, perhaps even larger than the sample size, maximum likelihood estimation might no longer be appropriate. The use of regularized estimators, such as ridge regression, lasso, scad, etc. requires adjustment to the FIC formulae. Even when the regularization takes automatic care of selection, Claeskens (2012) showed that selection via FIC is advantageous to get better estimators of the focus. Pircalabelu et al. (2016) used FIC for high-dimensional graphical models. For models with a diverging number of parameters FIC formulae using a so-called desparsified estimator have been obtained by Gueuning and Claeskens (2018). FIC may also be used to select tuning parameters for ridge regression. The focused ridge procedure of Hellton and Hjort (2018) is applicable to both the low and high-dimensional case and has been illustrated in linear and logistic regression models.

7. *Extensions to yet other models.* The methods explicated in section 3.1, yielding FIC machinery under a fixed wide model, can be extended to other important classes of models. The essential assumptions are those related to smooth log-likelihood functions and approximate normality for maximum likelihood estimators for the candidate models. Sometimes developing such FIC methods would take considerable extra efforts, though, as exemplified by our treatment in section 3.3 of linear mixed effects models. In particular, the methodology extends to models with dependence, as for time series and Markov chains with covariates (see Haug, 2019). This involves certain lengthier efforts regarding deriving expressions and estimation methods for the $K_{M,n}$ and $C_{M,n}$ matrices of (6, 7). Analogous FIC methods for time series are shown at work in Hermansen et al. (2016) for certain applications in fisheries sciences. Similar remarks also

apply to the advanced Ornstein–Uhlenbeck process models used in Reitan et al. (2012) for modeling complex layered long-term evolutionary data. Specifically, these authors studied cell size evolution over 57 million years, and entertained 710 candidate models of this sort. An extension of our paper's FIC methods to their process models is possible and would lead to additional insights in their data.

A challenge of a different sort is to develop FIC methods also when the models used are too complicated for log-likelihood analyses, but where different estimation methods may be used. A case in point are models used in Dennis and Taper (1994), for dynamically evolving times series models of the form $y_{t+1} = y_t + a + b \exp(y_t) + \sigma z_t$, met in density dependence analyses for ecology. These models do not have stationary distributions and special estimation methods are needed to analyse the candidate models.

DATA AVAILABILITY STATEMENT

One of the datasets in this manuscript is not publicly available. The access to the minke whale dataset is controlled by the scientific committee of the International Whaling Commission. Requests to access the datasets should be directed to IWC Scientific Committee's Data Availability Group (DAG).

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

ACKNOWLEDGMENTS

Some of our FIC and AFIC calculations have been carried out with the help of the R library `fic`, developed by Christopher Jackson, see `github chjackson/fic`, also available on CRAN; some of our algorithms are extensions of his. CC and NH thank Kenji Konishi and the other scientists at the Institute of Cetacean Research for obtaining the body condition data for the Antarctic minke whale, the IWC Scientific Committee's Data Availability Group (DAG) for facilitating the access to these data, and Lars Walløe for valuable discussions regarding the modeling. GC acknowledges support of KU Leuven grant GOA/12/14, and CC and NH acknowledge partial support from the Norwegian Research Council through the FocuStat project 235116 at the Department of Mathematics, University of Oslo. Finally and crucially, the authors express gratitude to three reviewers, for their detailed suggestions, which led to a better and more clearly structured article.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fevo.2019.00415/full#supplementary-material>

REFERENCES

- Berk, R., Brown, L., Buja, A., Zhang, K., and Zhao, L. (2013). Valid post-selection inference. *Ann. Stat.* 41, 802–837. doi: 10.1214/12-AOS1077
- Breiman, L. (2001). Statistical modeling: the two cultures [with discussion contributions and a rejoinder]. *Stat. Sci.* 16, 199–215. doi: 10.1214/ss/1009213726
- Burnham, K. P., and Anderson, D. R. (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, 2nd Edn.* New York, NY: Springer-Verlag.
- Charkhi, A., and Claeskens, G. (2018). Asymptotic post-selection inference for the Akaike information criterion. *Biometrika* 105, 645–664. doi: 10.1093/biomet/asv018
- Claeskens, G. (2012). Focused estimation and model averaging with penalization methods: an overview. *Stat. Neerland.* 66, 272–287. doi: 10.1111/j.1467-9574.2012.00514.x
- Claeskens, G., and Hjort, N. L. (2003). The focused information criterion [with discussion contributions and a rejoinder]. *J. Am. Stat. Assoc.* 98, 900–916. doi: 10.1198/016214503000000819
- Claeskens, G., and Hjort, N. L. (2008a). Minimizing average risk in regression. *Econometr. Theor.* 24, 493–527. doi: 10.1017/S0266466608080201
- Claeskens, G., and Hjort, N. L. (2008b). *Model Selection and Model Averaging.* Cambridge: Cambridge University Press.
- Claeskens, G., Magnus, J. R., Vasnev, A. L., and Wang, W. (2016). The forecast combination puzzle: a simple theoretical explanation. *Int. J. Forecast.* 32, 754–762. doi: 10.1016/j.ijforecast.2015.12.005
- Cunen, C., Hjort, N. L., and Nygard, H. M. (2019). Statistical sightings of better angels: analysing the distribution of battle deaths in interstate conflict over time. *J. Peace Res.* (accepted).
- Demidenko, E. (2013). *Mixed Models: Theory and Applications With R.* New York, NY: Wiley.
- Dennis, B., and Taper, M. L. (1994). Density dependence in time series observations of natural populations: estimation and testing. *Ecol. Monogr.* 64, 205–224. doi: 10.2307/2937041
- Dormann, C. F., Calabrese, J. M., Guillera-Aroita, G., Matechou, E., Bahn, V., Bartoň, K., et al. (2018). Model averaging in ecology: a review of Bayesian, information-theoretic, and tactical approaches for predictive inference. *Ecol. Monogr.* 88, 485–504. doi: 10.1002/ecm.1309
- Efron, B. (2014). Estimation and accuracy after model selection [with discussion contributions and a rejoinder]. *J. Am. Stat. Assoc.* 110, 991–1007. doi: 10.1080/01621459.2013.823775
- Gueuning, T., and Claeskens, G. (2018). A high-dimensional focused information criterion. *Scand. J. Stat.* 45, 34–61. doi: 10.1111/sjos.12285
- Hand, D. J., Daly, F., Lunn, A., McConway, K. J., and Ostrowski, E. (1994). *A Handbook of Small Data Sets.* London: Chapman & Hall.
- Haug, J. (2019). *Focused model selection criteria for Markov chain models, with applications to armed conflict data* (Master Thesis). Technical report, Department of Mathematics, University of Oslo, Oslo, Norway.
- Hellton, K. H., and Hjort, N. L. (2018). Fridge: focused fine-tuning of ridge regression for personalized predictions. *Stat. Med.* 37, 1290–1303. doi: 10.1002/sim.7576
- Hermansen, G. H., Hjort, N. L., and Kjesbu, O. S. (2016). Recent advances in statistical methodology applied to the Hjort liver index time series (1859–2012) and associated influential factors. *Can. J. Fish. Aquat. Sci.* 73, 279–295. doi: 10.1139/cjfas-2015-0086
- Hjort, N. L. (2014). Discussion of Efron's "Estimation and accuracy after model selection." *J. Am. Stat. Assoc.* 110, 1017–1020. doi: 10.1080/01621459.2014.923315
- Hjort, N. L., and Claeskens, G. (2003). Frequentist model average estimators [with discussion and a rejoinder]. *J. Am. Stat. Assoc.* 98, 879–899. doi: 10.1198/016214503000000828
- Jullum, M., and Hjort, N. L. (2017). Parametric of nonparametric: the FIC approach. *Stat. Sin.* 27, 951–981. doi: 10.5705/ss.202015.0364
- Jullum, M., and Hjort, N. L. (2019). What price semiparametric Cox regression? *Lifetime Data Anal.* 25, 406–438. doi: 10.1007/s10985-018-9450-7
- Ko, V., Hjort, N. L., and Hobæk Haff, I. (2019). Focused information criteria for copulae. *Scand. J. Stat.* doi: 10.1111/sjos.12387
- Liang, H., Zou, G., Wan, A. T. K., and Zhang, X. (2011). Optimal weight choice for frequentist model average estimators. *J. Am. Stat. Assoc.* 106, 1053–1066. doi: 10.1198/jasa.2011.tm09478
- Piracabelu, E., Claeskens, G., Jahfari, S., and Waldorp, L. (2016). A focused information criterion for graphical models in fMRI connectivity with high-dimensional data. *Ann. Appl. Stat.* 9, 2179–2214. doi: 10.1214/15-AOAS882
- Reitan, T., Schweder, T., and Hendericks, J. (2012). Phenotypic evolution studied by layered stochastic differential equations. *Ann. Appl. Stat.* 6, 1531–1551. doi: 10.1214/12-AOAS559
- Schweder, T., and Hjort, N. L. (2016). *Confidence, Likelihood, Probability: Statistical Inference with Confidence Distributions.* Cambridge: Cambridge University Press.
- Shmueli, G. (2010). To explain or to predict? *Stat. Sci.* 25, 289–310. doi: 10.1214/10-STS330
- Taper, M. L., Staples, D. F., and Shepard, B. S. (2008). Model structure adequacy analysis: selecting models on the basis of their ability to answer scientific questions. *Synthese* 163, 357–370. doi: 10.1007/s11229-007-9299-x
- Tibshirani, R. J., Rinaldo, A., Tibshirani, R., and Wasserman, L. (2018). Uniform asymptotic inference and the bootstrap after model selection. *Ann. Stat.* 46, 1255–1287. doi: 10.1214/17-AOS1584
- Tibshirani, R. J., Taylor, J., Lockhart, R., and Tibshirani, R. (2016). Exact post-selection inference for sequential regression procedures. *J. Am. Stat. Assoc.* 111, 600–620. doi: 10.1080/01621459.2015.1108848

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The reviewer, CJ, declared a past collaboration with one of the authors, GC, to the handling editor.

Copyright © 2019 Claeskens, Cunen and Hjort. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.