

Adaptive Harvest Management for the Svalbard Population of Pink-Footed Geese

**A Response to a Review of the:
*Progress Summary***

by F. A. Johnson, G. H. Jensen, and J. Madsen (February 1, 2013)

Conducted by:

Hans Baveco¹, Paul Goedhart² and Dick Melman²

On behalf of the Dutch Delegation to the

AEWA Svalbard Pink-footed Goose International Working Group



Prepared By:

Dr. Fred A. Johnson

Southeast Ecological Science Center, U.S. Geological Survey
Gainesville, Florida, USA

Gitte H. Jensen and Dr. Jesper Madsen

Department of Bioscience, Aarhus University
Aarhus, Denmark

May 21, 2013

¹ Alterra, Wageningen-UR, Hans.Baveco@wur.nl; Dick.Melman@wur.nl

² PPO/PRI Biometris, Wageningen-UR, Paul.Goedhart@wur.nl

Herein we provide excerpts of the reviewers' comments and our response. We did not duplicate the portion of the reviewers' comments in which they summarized our findings.

Reviewers' Comments:

First of all, we conclude that the authors are developing a very promising approach to determine harvest rates for use in adaptive management of the Pink-Footed geese, based on harvest maximization. They succeed in striking a good balance between simplicity and realism of models, as well as in exploiting all of the empirical data sources.

The report provides information on 3 aspects of the adaptive harvest management for Pink-Footed geese, a) the selection of models with their structural definition, the determination of model coefficients and their goodness of fit on the monitoring data; b) the calculation of optimal harvest rates for each possible state of the system (population and important environmental variables); c) the application and integration of the calculated optimal harvest rates in an adaptive harvest management cycle of e.g., three years.

Authors' Response:

We would like to express our sincere appreciation to the reviewers for the very thorough and constructive review of our Progress Summary. We emphasize that peer review is an essential aspect of the process of developing and implementing an adaptive harvest management (AHM) program for pink-footed geese, and we will continue to solicit reviews by the International Working Group and their staff, as well as scientists not directly engaged in this effort. We will also endeavor to publish science relevant to this effort in refereed professional journals, and we will make these documents available as soon as they are accepted for publication. In addition to reports describing development of the AHM process, an annual assessment will be produced each summer for the purpose of soliciting feedback and for guiding implementation of AHM. All project documentation will be available on the public AEWA website (<http://www.unep-awea.org/>).

Reviewers' Comments:

The authors are bound to use a two stage population model, because their statistical methodology does not allow to estimate parameters of a three stage model. We suggest a possible alternative way (state-space modeling, see Schaub and Abadi, 2011; Buckland et al., 2004) to estimate model coefficients for each of the 9 models discussed in the report, and thus enabling the use a three stage model. The state-space modeling approach would allow estimation of natural survival rates and avoid the a-priori assumption that harvest mortality amounts to one-half of total mortality. It also provides a structured way to incorporate additional data becoming available with continued monitoring. Also the capture-recapture based data used in the report, could be incorporated in the suggested approach.

Authors' Response:

This is an excellent suggestion and we have begun to explore opportunities for conducting this research. A comprehensive analysis of all capture-recapture data by Aarhus University is underway, and we have had initial conversations about how to integrate these data with other monitoring information. We note that there were several reasons we elected not to use state-space modeling at the outset of this project. The first was a time constraint, as it was necessary to develop some preliminary models quickly to help inform important discussions about management objectives and alternative harvest actions. Another reason is that initially we had some reservations about the accuracy of the harvest data, which we believed might be positively biased, and also because harvest data from both Norway and Denmark were not available for the full period of record. Finally, we note (as do the reviewers) that there is no information in the extant data that would allow us to estimate the size of the three age classes each year (unless you assume homogeneous natural mortality rates) or to determine whether they have different survival rates. Such heterogeneity in survival rates would be necessary to affect the optimal harvest strategy. We are hopeful that the new survival-rate analysis will determine whether such heterogeneity exists. Finally, we note that the way our models predict reproductive success is not dependent on an ability to separately estimate the number of sub-adults and adults.

Reviewers' Comments:

For the determination of optimal harvest rates we trust that the implementation of the method (stochastic dynamic programming) is correct (in theory, this could be checked from the source-code, by a person with hands-on experience with the program).

Authors' Response:

We appreciate the need to check the source code and have done that to the best of our ability (both the first and second authors have extensive experience with the software used for the optimizations). This summer, we will be attempting to derive optimal policies using different (and more computationally efficient) software and, thus, there will be two sources of optimal policies to compare, which will help verify the coding.

Reviewers' Comments:

From the text in the report we find it hard to grasp the exact meaning/definition of optimal harvest rates (see some of our comments to section V). A precise definition in words would help, in particular to understand the implications the choice of specific options/settings in the optimization method will have for application of the results in an adaptive harvest management cycle.

Authors' Response:

Harvest rate is defined as the proportion of the sub-adult and adult population alive at the beginning of the hunting season that is shot by hunters (both retrieved harvest and crippling loss). The harvest rate of young of the year is assumed to be twice that of older birds. As indicated in the report, the ability to regulate harvest rate is largely unknown, so we used a set of discrete rates that

spanned the range from no harvest to a rate that would initially be needed to reduce the size of the population. Since preparation of the Progress Summary, we have calculated optimal policies that prescribe an absolute harvest rather than a harvest rate.

Reviewers' Comments:

Our current understanding is that optimal harvest rate depends on current state of the system (population and environment) and is the rate of harvesting that maximizes – when exerted over a long period (infinity?) with stochastic fluctuations in demographic rates and some environmental variables – harvest under the condition that the total population stays around a target level (e.g. 60k with standard deviation 10k, which seems rather large to us).

Authors' Response:

The reviewers are correct that the optimal harvest rate depends on the current state of the system; i.e., the harvest policy is state-dependent. If the system state this year is different than last, a different harvest rate is likely to be optimal. In other words, the optimal policy prescribes a harvest rate that is appropriate for each state of the system that a manager might encounter at any time in the future. So while the harvest policy is constant, the optimal harvest rate is not. The key advantage of a harvest policy derived with stochastic dynamic programming is that it provides optimal harvest rates for all possible system states, absent any assumptions about equilibrium conditions or the system evolving as expected.

Reviewers' Comments:

Regarding the weighting of the models, the text suggests that weights are adjusted during the optimization process. On the other hand, it is stated everywhere that results apply to the situation with equal weights assigned to models. This brings us to the more theoretical question whether it is more clear/insightful to calculate a weighted average harvest rate over 9 models in every step of the optimization program, or to compare optimal harvest rates calculated for each model separately. As an alternative an uncertainty analysis using a simple single model, e.g. with a single survival and reproduction parameter and no density dependence, could be performed.

Authors' Response:

The process we are proposing is termed passive adaptive management, in which the updating of model weights occurs “off-line” from the optimization process. This means that possible changes in model weights are not anticipated in the optimization. Simply put, this means we derive an optimal policy based on current understanding of system dynamics (as captured in the model weights). An optimal action is chosen from the policy and then each model predicts a different system state in the following year. When new monitoring information is available, predicted and observed system responses are compared and the model weights updated. Then a new harvest policy is derived using these new weights. This process is repeated indefinitely. The use of equal model weights throughout the report reflects our belief that the model-weight updating depicted in Fig.12 is unreliable, absent reliable estimates of actual harvest rates over the period of record.

Regarding the use of multiple models and weights, we recognize that calculation of model-specific harvest policies is possible, but it would not be appropriate to simply “average” the policies. Passive adaptive management based on stochastic dynamic programming provides a formal and coherent approach for coping with model uncertainty. Also, we note that in a subsequent analysis, we determined that a harvest policy derived using equal model weights is robust to model uncertainty; i.e., it minimizes the maximum loss in management performance that could arise from not knowing the most appropriate model of system dynamics.

Reviewers’ Comments:

Missing from the report is a proposal for the way to integrate optimal harvest rates in a three-year management cycle. In the end, it is this aspect that will allow us to judge the usefulness of the optimization approach, and to evaluate whether its clear advantages outweigh its drawbacks (complexity of the methodology). To us, it is not immediately clear whether this is the case. Optimal harvest rates given a current state of the system are (if understood correctly) based on a kind of expected/average/uncertain future (accounting for environmental stochasticity like the number of warm days in May and for random variation in demographic rates). The realization in the three years of the management cycle following the choice of an optimal harvest rate will always be a particular sequence, of e.g., 2 warm May months and 1 cold. In hindsight, the adopted management action (a particular harvest rate) will not be optimal at all. Of course, AHM is adaptive, and designed to handle this situation by adjusting the harvest rate to one that is considered optimal for the updated system state, and also by updating the weight given to each model (if more than one) by a new fit of models on the extended set of monitoring data. But it also begs the question whether the proposed harvest rate really needs to be the long term optimal one and cannot be replaced by a harvest rate obtained from a simpler, non-optimizing modeling approach, e.g., 3 year stochastic simulations projecting population size and structure for a range of harvest rates, with each of the nine models.

Authors’ Response:

We have now computed preliminary policies of optimal harvest quotas for 3-year periods. We still use stochastic dynamic programming, but now the harvest policy prescribes a constant harvest quota based on the system state at the time of the decision and the expected evolution of the system over three years. So far, this expectation of change in system state is deterministic because computation of a fully stochastic policy is computationally intractable with our current software. This is the principal reason we will be exploring new software this summer.

Again, we stress that the harvest policy is state dependent and does not provide a harvest rate (or harvest) that is the “long term optimal one.” However, the reviewers are correct to imply that there will be a loss in management performance by giving up the option of making harvest decisions annually. We recently completed extensive optimization and simulation exercises that demonstrated: (a) there is generally more variability in population size under a 3-year cycle of decision making than under annual decisions; (b) there is greater risk of low populations (i.e., $N \ll 60k$) with a 3-year cycle than with a 1-year cycle; (c) there is considerable risk of the population escaping our ability to control it (with a maximum harvest quota = 20k) if the most appropriate

model lacks density dependence; and (d) the magnitude of change in harvest quota is substantially higher in a 3-year cycle than in a 1-year cycle (i.e., relatively infrequent changes in harvest quota are accompanied by relatively large changes in the quota when a change is necessary).

Reviewers' Comments:

In general it would be interesting to get a feeling for the added value of the optimization approach. Comparison of results for the nine models in a stable state at 70k (discussed in the upper part of page 14) with Monte Carlo simulation results of optimal harvest rates (Table 2 on page 19) could be enlightening in this respect. However, a systematic comparison is not possible based on the information provided in the report.

Authors' Response:

We are unclear what is being suggested. The nine models do not suggest stable states but rather stability regimes under a given harvest policy (quasi-stationarity), and it is well known from harvest theory that constant rates of harvest will under-perform state-dependent rates in a stochastic system. And, unless we misunderstand the reviewers, we don't understand why it would ever be preferable to use a sub-optimal approach when an optimal one is available.

Reviewers' Comments:

The importance/relevance of applying an optimization approach could be tested by comparing the results for harvest maximization, as described in the report, with results for harvest minimization. This would give us direct information on the range of outcomes using an optimization approach. When the bandwidth between maximized and minimized harvest is small, optimization is not really worth the trouble. In addition it would make the methodology useful for the management of e.g., more numerous species, like the steadily increasing breeding population of Greylag geese in the Netherlands, where managing the population with minimal hunting effort appears to be more of an issue.

Authors' Response:

We are unsure whether the reviewers are questioning the application of optimization methods or the appropriateness of a management objective that seeks to *maximize* sustainable harvest, subject to a population constraint. If it is the latter, we initially considered an approach in which the objective was to *minimize* harvest, subject to the population constraint. The resulting policy was deemed unacceptable because implementation would be characterized by several years of no hunting followed by a year in which a massive harvest was necessary. We also explored an objective solely focused on achieving the population target. The resulting policy was also deemed unacceptable because it produced extreme variability in harvest quotas. So we believe that the current objective function represents a good balance of concerns, recognizing that there is considerable room for debate about what is the desirable range of population sizes (assuming that the target of 60k could be achieved on the average).

Reviewers' Comments:

Page 5. For survival the assumption that harvest mortality is one-half of total mortality seems quite drastic. A first check could be to compare the harvest numbers in DK (1990-2002) (assuming no harvest in N) to estimated survival of the monitored population (HarvDen, Survival and censusN in the table of Appendix 1). The results indicate that at the first half of this period the (minimum) estimate of harvest mortality is higher than one-half of total mortality, while in the second half it is approximately equal to one-half. Is this check too simple, e.g. because harvest estimates are too unreliable for this period?

Year	censusN	Annual Survival	Annual Mortality	Harvest Denmark	Natural Mortality	Fraction Natural Mortality
1990	26000	0.9201	2077	1800	277	0.011
1991	32500	0.9672	1066	3000	-1934	-0.060
1992	32000	0.9321	2173	2500	-327	-0.010
1993	34000	0.9328	2285	2300	-15	0.000
1994	33000	0.9256	2455	2600	-145	-0.004
1995	35000	0.9171	2901	2800	101	0.003
1996	33000	0.8531	4848	2000	2848	0.086
1997	37500	0.8917	4061	2500	1561	0.042
1998	44800	0.9149	3812	1414	2398	0.054
1999	38500	0.9315	2637	1973	664	0.017
2000	43100	0.8639	5866	2567	3299	0.077
2001	45000	0.8671	5981	2353	3628	0.081
2002	42000	0.8743	5279	2611	2668	0.064

It seems quite unfortunate that no age-dependent survival estimates are available. In our appendix A we show that with a state-space model, all available data (in appendix 1) can be used to estimate survival for sub-adults and adults separately, assuming the 3-stage model described in the next section.

Authors' Response:

Unfortunately, checking the assumption in this manner is not as straightforward as it may seem. First, recall that the survival rate estimates apply to the period Feb 1 in year t to Jan 31 in year $t+1$, while the population census is on about Nov 1 in year t . So the first step would be to align the anniversary date of the survival rate with Nov 1, which requires an assumption about the distribution of total mortality throughout the year. While we assumed elsewhere that natural mortality is evenly distributed throughout the year, it does not seem reasonable to assume that total mortality is evenly distributed because of the harvest mortality concentrated in the autumn. Moreover, total harvest in any given year includes the harvest of surviving individuals from the previous year, as well as harvest of young-of-the-year. If we could reduce the total harvest by the harvest of young, we could compare the harvest of survivors from the previous year with the total mortality of survivors and arrive at an estimate of the proportion of total mortality attributable to hunting. Unfortunately, we currently have no way to partition the annual harvest into young and older birds, but we do know that young are about twice as vulnerable to harvest as older birds so they make up a disproportionately large portion of the total harvest. For this reason, it can appear that harvest mortality is a larger portion of total mortality than is actually the case. In conclusion, limitations of the data prevented us from rigorously checking our assumption about the proportion of total mortality attributable to hunting, but we believe it is acceptable for the reasons provided in the report. We also note that the assumption was made only for the period 1990-2002, as we believe the dramatic increase in harvests in Denmark and Norway in more recent years could have well changed the proportion of mortality due to hunting (as well as the annual survival rates themselves).

We agree it would be useful to have age-dependent estimates of survival and we hope that our recently initiated analysis will provide these. As the reviewers note elsewhere (see Appendix A), however, a state-space model using data available in the report cannot be used to estimate age-dependent survival (unless sub-adults can be distinguished from adults in the population census).

Reviewers' Comments:

Page 6. TempDays and TempSum are considered as the only weather covariates. The correlation between these two variables equals 0.86. It is therefore not surprising that TempDays and TempSum do equally well in regression models for survival and reproduction.

Authors' Response:

Agreed.

Reviewers' Comments:

Page 8. The 3-stage model is used only to make one-year ahead predictions for the period 1990-2002. We were able to reproduce Figure 3 using the data in Appendix 1. The table below shows the quantities used; the first 4 columns are taken from Appendix 1, and the remaining columns are calculated according to the formulae at the bottom of page 7 and the top of page 8.

Year	censusN	censusNjuv	Survival	θ_t	$\theta_t^{0.25}\theta_{t+1}^{0.75}$	$1-h_t$	p_t	\widehat{N}_t
1990	26000	3224	0.9201	0.9584	0.9770	0.9601	0.1240	-
1991	32500	7215	0.9672	0.9833	0.9694	0.9836	0.2220	32116
1992	32000	1984	0.9321	0.9649	0.9651	0.9661	0.0620	32449
1993	34000	6154	0.9328	0.9652	0.9623	0.9664	0.1810	36443
1994	33000	4092	0.9256	0.9614	0.9579	0.9628	0.1240	35961
1995	35000	8260	0.9171	0.9568	0.9296	0.9586	0.2360	39661
1996	33000	6072	0.8531	0.9207	0.9372	0.9265	0.1840	36944
1997	37500	5400	0.8917	0.9427	0.9523	0.9459	0.1440	34174
1998	44800	5466	0.9149	0.9556	0.9623	0.9575	0.1220	38945
1999	38500	4736	0.9315	0.9645	0.9362	0.9657	0.1230	47474
2000	43100	2112	0.8639	0.9270	0.9284	0.9320	0.0490	35323
2001	45000	4905	0.8671	0.9288	0.9319	0.9335	0.1090	41923
2002	42000	4452	0.8743	0.9329	-	0.9371	0.1060	43960

These values can be used to check whether the assumption holds that hunting mortality is one-half of total annual mortality. The assumed amount of hunting equals $h_{t+1}N_t$ and this can be compared with the sum of HarvDen and HarvNor taken from Appendix 1. In doing so, the missing values for HarvNor are replaced by 250. It appears that the assumption does not hold too good for the years 1991-1997 and for the year 1999 (compare last two columns of the table below). This could imply that natural survival is larger than assumed in this report.

Year	censusN	HarvDen	HarvNor	h_t	HarvDen+HarvNo	$h_{t+1}N_t$
1990	26000	1800	250	0.0399	2050	-
1991	32500	3000	250	0.0164	3250	426
1992	32000	2500	250	0.0339	2750	1103
1993	34000	2300	250	0.0336	2550	1075
1994	33000	2600	250	0.0372	2850	1265
1995	35000	2800	250	0.0414	3050	1368
1996	33000	2000	250	0.0735	2250	2571
1997	37500	2500	250	0.0541	2750	1787
1998	44800	1414	250	0.0425	1664	1596
1999	38500	1973	250	0.0343	2223	1534
2000	43100	2567	250	0.0680	2817	2620
2001	45000	2353	400	0.0665	2753	2864
2002	42000	2611	500	0.0629	3111	2828

Authors' Response:

See our previous comments regarding the partitioning of harvest rates and natural mortality.

Reviewers' Comments:

The authors state that Figure 3 suggests “that survival and reproductive estimates were unbiased”. Although stated cautiously this is a rather strong statement.

Note also that different partitions of Survival into natural mortality and harvest rate may all more or less reproduce Figure 3. The formulae on top of page 8 all have the same multiplication factor $\theta_t^{0.25} \theta_{t+1}^{0.75} (1 - h_{t+1})$. This factor is more or less the same for every partition of Survival into natural mortality and harvest rate. Taking different values for this partition, i.e. different values of δ in $h_t = (1 - S_t)/\delta$, the multiplication factors below are obtained. Note that small values of δ imply large hunting rates, and large values of δ imply small hunting rates.

Year	$\delta = 1.1$	$\delta = 1.5$	$\delta = 2$	$\delta = 4$	$\delta = 8$
1990	0.966	0.963	0.961	0.958	0.957
1991	0.933	0.935	0.937	0.939	0.940
1992	0.933	0.933	0.933	0.933	0.933
1993	0.926	0.926	0.927	0.927	0.927
1994	0.917	0.918	0.918	0.919	0.919
1995	0.855	0.859	0.861	0.865	0.867
1996	0.891	0.888	0.886	0.884	0.883
1997	0.914	0.913	0.912	0.910	0.910
1998	0.931	0.930	0.929	0.928	0.928
1999	0.866	0.870	0.873	0.877	0.878
2000	0.867	0.867	0.867	0.866	0.866
2001	0.874	0.874	0.873	0.873	0.873
2002	-	-	-	-	-

This implies that something very similar to Figure 3 can be obtained for a large range of δ values.

Authors' Response:

It was not our intent to imply that the analysis in questioned supported our assumption that harvest mortality was half of total mortality. Rather we were interested in whether estimates of *annual* survival (including all mortality) and reproductive rates implied rates of change in population size that were consistent with those observed in the population census. We felt obliged to perform this check because of evidence that estimates of survival and reproductive rates of North American waterbirds imply rates of population growth much higher than those observed in population counts. Fortunately, rates of population growth from the pink-foot census and those based on estimates of survival and reproductive rates were similar (i.e., regression coefficients of about zero and one).

Reviewers' Comments:

Page 8. JJM discard the 3-stage model and replace it by a 2-stage model, mainly because there are no survival estimates and no census data available distinguishing between sub-adults and adults. JJM state that "Pink-footed geese may not be sexually mature until age three". Assuming sub-adults

and adults to reproduce equally, thus seems like an undesirable simplification, leading to underestimates of per-capita reproduction. JJM also assume that “Hunting mortality was additive to natural mortality and a constant one half of total annual mortality”. This is an important assumption since, see page 21, “survival is the most critical rate determining an appropriate harvest strategy”, that should – if possible – not be made a-priori. We show that it is possible to fit a three stage model, to disentangle natural mortality from mortality through hunting and to estimate survival, using the whole time-series of census and harvest (appendix A).

Authors’ Response:

We do not assume that sub-adults and adults produce equally. Rather, we use the pooled number of sub-adults and adults in year t to *predict* the observed number of young in year $t + 1$; i.e., we merely take advantage of a statistical correlation, absent any assumption about the underlying mechanism. See our previous comments about the partitioning of annual mortality.

Reviewers’ Comments:

Page 9. The authors state that “the behavior of models outside the range of experience is often more important than that for which data are available”. This is the rationale for fitting different models for both survival and reproduction. An alternative could be to use a single model but with different parameter values which cover the range of plausible values.

Authors’ Response:

Agreed. A Bayesian state-space model could be used, in which possible (and properly correlated) values of the parameters are drawn from the joint posterior distribution. As mentioned, we intend to explore the feasibility of this approach.

Reviewers’ Comments:

Page 10. Survival model (1). All survival models were fitted using the important assumption that hunting mortality is one-half of total annual mortality. Moreover these models are fitted using data in the years 1990-2002 only. This is a period in which the population is growing slowly, in contrast to the larger growth in later years. The first model employs a Beta distribution. We were able to reproduce the parameter estimates 125.16 and 6.46 (we used maximum likelihood instead of the method of moments to obtain estimates 124.50 and 6.42). The Beta distribution is then discretized to a coarse grid. This seems unnecessary since drawing from a Beta distribution is rather simple.

Authors’ Response:

See our previous comments about the partitioning of annual mortality. The distributions of continuous random variables were approximated by discrete variables because the backward-recursion algorithm of stochastic dynamic programming requires all states, actions, and sources of variation to be discrete.

Reviewers' Comments:

Page 10. Survival model (2). We were almost able to reproduce the parameter estimates at the top of page 11. Using data for the years 1990-2002 we obtained $\text{logit}(\theta) = 2.770 + 0.0495 X_t$

Authors' Response:

The reviewers' parameter estimates would differ slightly from ours if they used the survival rates from Appendix 1 of the report because they have an anniversary date of February 1, rather than the estimates we used, which have an anniversary date of November 1 (see page 7 of the report).

Reviewers' Comments:

Page 11. Survival model (3). Ordinary regression is applied to transformed observations. This assumes that the variance of the transformed values is constant. It is better to use logistic regression to fit this model. We were unable to reproduce the parameter estimates at the bottom of page 10. Using data for the years 1990-2002 we obtained $\text{logit}(\theta) = 4.114 + 0.0488 X_t - 0.0366 N_t$. JJM state about this model that it "involves extrapolating beyond the limits of the data and thus lacks empirical evidence". It may be hazardous to apply this model to later years with larger population sizes, since there does not appear to be a strong density dependence in the data (see Appendix 1). Of course possible density dependence in later years could be masked by warmer springs.

Authors' Response:

See our comment above concerning the possible reason for slight discrepancies in parameter estimates. We agree that a logistic regression for survival would be more appropriate, but we did not have access to the raw mark-recapture data at the time of this preliminary work. However, we could not reject the null hypotheses of constant variance and normality of errors ($P > 0.5$) for any of the survival models.

Reviewers' Comments:

Page 11-12. Reproduction models. Data from 1980 onwards were used to fit the three models. We obtained approximately the same parameter estimates for all three models. The last model, i.e. the beta binomial model with no covariates was again discretized. Again this seems unnecessary.

Authors' Response:

See above comment about the requirement to use discrete random variables in stochastic dynamic programming.

Reviewers' Comments:

Page 12. We were able to reproduce the estimates of the beta-binomial model for temperature days. Discretization of this distribution seems unnecessary.

Authors' Response:

See above response.

Reviewers' Comments:

Page 13. The three models for survival and reproduction are then combined into 9 different models. Using the equation at the top of page 14 (which can be derived from the 2 equations at the bottom of page 8) the required harvest rate can be calculated. Using $Y=15.4k$, $A=54.6$, $X=10$, and applying the fitted models we obtained the values in the table below.

model	θ_i	p_i	R_i	h_i^*	Harvest $h_i^* \times 70k$
M0	0.9510	0.1222	0.1393	0.0770	5391
M1	0.9619	0.1222	0.1393	0.0874	6121
M2	0.8511	0.1222	0.1393	-0.0314	-2196
M3	0.9510	0.1520	0.1792	0.1083	7581
M4	0.9619	0.1520	0.1792	0.1184	8287
M5	0.8511	0.1520	0.1792	0.0036	252
M6	0.9510	0.1400	0.1628	0.0957	6698
M7	0.9619	0.1400	0.1628	0.1059	7414
M8	0.8511	0.1400	0.1628	-0.0105	-735

Note that the models do not differ much in their value of R_i , but do differ considerably in their value of θ_i . Especially the value for the density dependent survival model does seem to produce a rather low value of natural survival. For the observed population size in 2011 (80k) the survival rate of this model is even as low as 0.786. This is due to extrapolation (which the authors noted on page 10).

The authors state that for models M6 and M7 a harvest of approximately 17k is required. We arrive at a required harvest of 7k for these models.

Authors' Response:

The discrepancy is because we used the observed reproductive rate for 2010 for models M6, M7, and M8 (0.195) rather than the model prediction of 0.140. The reviewers approach is more consistent with the predicted values used in the other models. We note this discrepancy underscores the need to consider uncertainty in predicted vital rates, as well as the desirability of state-dependent harvest decisions (i.e., those based on the observed rather than predicted state of the system).

Reviewers' Comments:

Page 14. The equation at the top of page 14 can be used to derive the required harvest rate to obtain a stable population of 60k under various parameter values θ and R . This will give a range of required harvest rates. These parameter values can be assigned prior probabilities by using data, models and/or expert opinion, resulting in a mean harvest rate along with a standard error. This simple and robust approach can be also used for a population of 80k to see what has to be done extra to arrive at a population of 60k.

Authors' Response:

Again, we emphasize that we are not seeking a mean harvest rate or even a harvest rate for any particular year. We seek a policy of state-dependent harvest rates, that prescribes an optimal harvest rate for every system state that might be observed in the future, and that will be optimal whether the future unfolds as expected or not (i.e., Bellman's Principle of Optimality).

Reviewers' Comments:

Page 14. The likelihood at the bottom of page 14 employs a model specific prediction of population size. This is, as we assume, given the observed population size in the previous year. If this is correct, the comparison of the models is based on the one-year ahead predictions. This a rather limited comparison of the models since the starting point is always "correct". The difference between the models becomes especially apparent after repeatedly applying the same model, as can be seen in the table at the bottom of page 13. Given this limited comparison the authors conclude that the 9 models should be assigned equal prior probability. However there seems scope to assign a (much) lower probability to the density dependent survival model.

Authors' Response:

The reviewers are correct that the likelihood is based on a comparison of one observation with one model-specific prediction. As described, multi-year comparisons are accommodated by the successive application of Bayes' theorem. Given that we have a 10-year gap in which model weights can be updated, we suggested that it might be appropriate to begin the adaptive-harvest management process with equal model weights (i.e., complete uncertainty as to which model is currently most appropriate). Moreover, an optimal policy derived using equal model weights is expected to be robust (i.e., will perform reasonably well regardless of the most appropriate model).

Nonetheless, we concede that there is room for debate about initial model weights in the absence of up-to-date empirical information.

Reviewers' Comments:

Page 16-18. Some aspects of the applied method (optimization, Stochastic Dynamic Programming) are not completely clear from the text. The description of the method suggests that model weights are updated *inside* the optimization procedures. However, it is also stated at many locations in the text that results apply to the case of equal weights assigned to the models. The optimal harvest rate is calculated for the period $t=0$ to $t=T$. Is T in this case infinity, or just a sufficiently large value, or related to a much shorter period? Which sources of stochasticity are accounted for in the optimization? The demographic rates and, for the models with weather variables, the number of temperature days in May? Or are there more?

Authors' Response:

As described previously, updating of model weights occurs “off-line” from the optimization process. This means that possible changes in model weights are not anticipated in the optimization. Simply put, this means we derive an optimal policy based on current understanding of system dynamics (as captured in the model weights). An optimal action is chosen from the policy and then each model predicts a different system state in the following year. When new monitoring information is available, predicted and observed system responses are compared and the model weights updated. Then a new harvest policy is derived using these new weights. This process is repeated indefinitely. The use of equal model weights throughout the report reflects our belief that the model-weight updating depicted in Fig.12 is unreliable absent reliable estimates of actual harvest rates over the period of record.

The backward-recursion process in the optimization proceeds over an indefinitely long time period until the state-dependent policy stops changing. This convergence to a time-independent policy is the appropriate policy to use for maximizing the objective value over an infinite time horizon. Sources of stochasticity considered in the optimization are variation in temperature days, random variation in survival (models M6, M7, and M8) and in reproduction (models M0, M3, and M6), as well as uncertainty about the models themselves (i.e., the model weights describe a discrete probability distribution of models).

Reviewers' Comments:

page 18. Harvest utility is given a mean value of 60k and a standard error of 10k. The standard error seems rather large. Does the value of the standard error affect the conclusions in any way?

Authors' Response:

A standard deviation of 10k was chosen to reflect managers' opinion that population sizes in the range of 50 – 70k were largely acceptable, given that we likely have limited ability to control population size against the background of environmental variation. The working group has since adopted a revised utility function, which imposes a more severe decline in utility beyond this range.

We have explored a number of utility functions that maintain relatively high utility for a population in the range 50 – 70k and the differences in the optimal harvest policies are minor. Nonetheless, we stress that the shape of the utility function is an inherently subjective (i.e., value-based) judgment on the part of managers, and that substantive changes can dramatically affect the optimal policy.

Reviewers' Comments:

page 19. The table shows that under models M2, M5 and M8 (the density dependent survival models) the population size is approx. 56k instead of 60k as required. Is this due to the discrete values for harvest rate which are used in the simulation?

Authors' Response:

No, the difference is due to the fact that the optimal policy seeks a compromise between maximizing harvest and attaining the population target (as expressed in the objective function on page 17).

Reviewers' Comments:

page 20. I would be very nice indeed to have more data on the age structure of the population, either in spring (after reproduction) or in the autumn. Moreover the age structure of the harvest is also valuable information. Note that in our alternative model there is less need to have a census at another time as long as (1) harvest can be partitioned into pre/post census and (2) it can be assumed that natural mortality mainly is in winter and after harvesting.

Authors' Response:

As mentioned previously, we agree that more information on the age structure of the population would be helpful and we have begun some new efforts in this regard. For now, like the reviewers, we assumed that all harvest occurs prior to the population census. Unlike the reviewers, however, we assumed that natural mortality was distributed evenly throughout the non-hunting season. Unless we can discern differences in survival rate among age classes, or we are able to distinguish sub-adults and adults in autumn, we see little advantage in the state-space modelling approach other than the ability to incorporate harvest data (which is a desirable feature).

Reviewers' Comments:

page 20. We agree with the authors that it is more suitable to maximize harvest instead of harvest rate, when following an optimization method. As the authors show in appendix 2, this requires a pre-harvest census instead of post-harvest census

Authors' Response:

The authors now have a prototype optimization framework, in which the control is absolute harvest rather than harvest rate, and in which the harvest quota is set only once every three years. We reported preliminary results from this approach at the April 2013 meeting of the International Working Group and expect to have a formal report later this summer.

Again, we would like to thank the reviewers for their time and substantive comments. We welcome this continuing level of scrutiny with both development and implementation of adaptive harvest management for the pink-footed goose.