The Meta-analysis of Population Trends of Loggerhead and Leatherback Turtles

1 Introduction

Critical conservation decisions often have to be made from widely scattered local data that may show what appears to be contradictory trends. Traditionally, each trend was examined independently to determine if they show statistically significant results. This approach—determining the proportion of studies individually exhibiting statistical significance in the prescribed direction—is known as "vote counting" in the meta-analysis literature, and is inherently flawed (Hedges and Olkin 1985). Hedges and Olkin (1980) showed that as the number of studies becomes large, the proportion of studies yielding significant results is approximately the average power of the test.

Here we suggest an improved approach. Our goal is to obtain the most powerful estimates, and still model the variation among sites.

2 A Mixed Model Approach

We begin with the simplest of models. Let X_{it} be the estimate of abundance of nesting females at each site i in year t. We will examine the simplest possible dynamic model for each site, given that each site began with an initial number of nesting females in the first year of the study, i.e. at $t = 0$ each site had x_{i0} nesting females. If each population is changing at a constant rate over the period of time of the censuses, we have

$$X_{it} = X_{i0}e^{r_i t + \varepsilon_{it}}$$

where r_i is the instantaneous rate of population change, X_{i0} is the initial population size of the ith population, and ε_{it} is the error in the estimate of abundance and deviations from the assumed model (sometimes called "process error").

It is unlikely that all the nest sites will have exactly the same rate of change, so we will investigate a simple mixed effect model where we assume that $r_i \sim N(\mu_r, \sigma_r^2)$, where μ_r is the mean instantaneous rate of change in population size for the nesting sites and σ_r^2 is the variance among populations.

We will investigate two approaches to the above model: (1) we log transform the data and use a linear mixed model and (2) we will use the raw counts and use a generalized linear mixed model. the second approach is more flexible, and can handle years in which no turtles were observed, but the first is easier to implement and understand.

3   A Linear Mixed Model

First, we can log transform the above model. Let $x_{it} = \log X_{it}$, then we have

$$x_{it} = x_{i0} + r_i t + \varepsilon_{it}.$$

Now r_i can be interpreted as a slope, and x_{i0} can be interpreted as an intercept. The initial population number is also observed with error so that it is clearer if we define the true initial log abundance to be a parameter, $\theta_i = $ true x_{i0}, which we will estimate. The above equation becomes

$$x_{it} = \theta_i + r_i t + \varepsilon_{it}.$$

The nesting sites vary in size and suitability for nesting, and these properties can be thought of as intrinsic to each site, and unrelated to any other site. Since the intercept for each nesting site determines the initial density, we treat the intercepts as site-specific fixed effects. The rate of change of the population at each nesting site, however, may reflect larger-scale phenomena such as climate or management policies. Therefore we model the slope for each nesting site as a random effect.

We used restricted maximum likelihood (REML) to fit the linear mixed models. REML can be thought of as an adjustment to the degrees of freedom accounting for the fixed effects, giving unbiased variance estimates (Searle et al. 1992). The likelihood ratio test (lrt) was used to compare the fit of different models.

We estimated the parameters of this model under two different assumptions for ε_{it}. First, we assumed that the variance of ε_{it} is the same for all sites. That is, we assume that $\varepsilon_{it} \sim N(0,\sigma^2)$, where σ^2 is estimated from the data. An alternative approach is to assume that the variances are primarily due to factors unrelated to sample size, and estimate a separate variance, σ_i for each site.

4   A Generalized Linear Mixed Model

In order to test the robustness of our approach we investigated alternative model formulations. Alternatively, we assumed that the residual variation in the above model was described by a gamma as opposed to a lognormal distribution. For this model, we used a generalized linear mixed model with a log link and a gamma error distribution. In this model, the variation of r among sites is still Gaussian. These parameters were estimated using the generalized linear mixed model methods developed by Wolfinger and O'Connell (1993).

5   Obtaining an overall estimate of population changes

The above model does not provide an estimate of the total population change over time. We suggest two approaches to obtain such a change; which require different amounts of data.

5.1 Trend data available for all sites

In this situation we have some data for all sites; however at some sites very limited data might be available. In this case we can include all reliable data and obtain BLUP for each site and year. For each site we obtain an estimate.

6 Data

6.1 Loggerheads

We used data from beach surveys of nesting. We have limited our analysis to beaches where we believe the effort has been relatively constant over time. However, this assumption may not always be true, particularly for the early years. We view this as the greatest uncertainty in the analysis.

The North Carolina beaches used in the analysis were: BBVO BHIC CHNS CLMB CLNS HBSP PINWR. We analyzed this group as a unit.

The Georgia beaches used in the analysis were: BKB CUM JEK LCI LSI OSS PIN SAP SCI SEA SSI TYB WAS WMI. For one of the Georgia beaches, LCI, there has been a large amount of beach erosion. We decided to keep this beach in the analysis because this represents natural variation among beaches, and the turtles from that beach probably nested elsewhere in the region..

We decided not to use the best studied beach in South Carolina, Cape Island, because predator control appears to have made this beach unrepresentative.

We combined results for two different time periods: 1989-1999 and 1979-1999. We believe that the data from the more recent period are more consistent.

7 Results

We fit a variety of models, to each data set: with lognormal, gamma, or extra-Poisson variability in the observation error, with separate error variances for each beach, and with and without outliers. We found that there was relativiely little difference among these models fits, but that the separate error variance was usually needed. For simplicity we will report the lognormal error with separate error varienaces as our primary estimates.

7.1 Loggerheads

7.1.1 Georgia

We first discus the period from 1979-1999.

The estimates of the slope if each nesting site is fit individually by OLS regression are:

```
Nesting     Parameter       Standard     T for H0:
Site        Estimate            Error    Parameter=0    Prob > |T|
BKB         0.005577     0.01366527          0.408         0.6878
```

3

CUM	0.007004	0.01254062	0.559	0.5860
JEK	0.015879	0.01160858	1.368	0.1903
LCI	-0.092834	0.00900179	-10.313	0.0001
LSI	0.041901	0.03072965	1.364	0.2026
OSS	0.017643	0.01747988	1.009	0.3270
PIN	-0.011962	0.03620422	-0.330	0.7473
SAP	0.057755	0.02783523	2.075	0.0622
SCI	-0.011230	0.01948592	-0.576	0.5760
SEA	0.073510	0.04804723	1.530	0.1604
SSI	-0.192613	0.06439019	-2.991	0.0152
TYB	-0.001411	0.08965388	-0.016	0.9878
WAS	0.024396	0.01453494	1.678	0.1096
WMI	-0.313002	0.42583946	-0.735	0.5965

Note that there are two nominally significant declines (LCI and SSI), but 8 of the 14 slopes are positive.

We now consider the mixed model results. The estimate for the mean instantaneous rate of change, μ, was -0.0060 (SE $= 0.0168$). The standard deviation of r among sites, σ_r, was estimated to be 0.05, and the standard deviation of the error, σ, was estimated to be 0.45.

If we allowed separate error variances for each site, μ is estimated to be slightly positive ($\hat{\mu} = 0.001$, SE $= 0.014$), and with more variation among sites ($\hat{\sigma}_r = 0.04$).

Similar results were obtained for the time period from 1989 to 1999 when the nesting data is believed to be more reliable, i.e. there was no evidence of an overall change in the populations in Georgia.

7.1.2   Northern US population

In this analysis we fit the data from 78-89 with one "r" and the data after with a second "r" using a piecewise linear model. In the results below, the r for the early period is estimated, then the change in "r" is added on, all in the same mixed model.

observational error assumption	separate error variance	time	r	Std Error r	DF
lognormal	no	> 1978	0.0052	0.017	29
lognormal	no	> 1989	0.023	0.01933172	22
lognormal	yes	> 1978	-0.00891	0.012	29
lognormal	yes	> 1989	0.039	0.015	22

The effect of the mixed model analysis for the Northern US leatherback populations is clear when examining the data (Fig. 1). Each point in the plot represents the number of log transformed nests. Note that when there are outliers, (Bogue Banks), the mixed model fits downwait these outliers.

Figure 1: Plots of log transformed nest counts for the populations in the northern US population of leatehrback turtles. Each point is an observation, the solid line is the ordinary least squares fit, the dashed line is the BLUP from the mixed model. The vertical dashed line is the time that TED's were employed.

The upshot is: no change from 78 to 89, and a probably increase from 89 to 1999. The magnitude and statistical significance of the increase depends upon the exact implementation of the model, but the effect is suggestive in all analyses.

7.1.3 South Florida Loggerheads

The results for the south Florida loggerheads are

observational error assumption	separate error variance	time	r	Std Error r
lognormal	no	> 1978	0.06220707	0.01378918
lognormal	no	> 1989	-0.02726639	0.02226328
lognormal	yes	> 1978	0.066	0.01323360
lognormal	yes	> 1989	-0.03255165	0.02042191

In summary, for south Florida the model with the separte error variances for each beach, the population has been growing with a $\hat{r} = 0.066$ (s.e. 0.013) since 1979. When an additional term is added for the period since 1989, there is not a signifianct change (added r is estimated to be 0.032 (se 0.02). Note that the direction of change is the opposit sign as for the northern population.

7.2 Leatherbacks

For leatherbacks we treated the data from the Virgin Islands, South America and Floria as separte groups, or populations.

In all cases, we examined data from 1979 on, for Virgin Isalns and Florida, but used since 1987 for South America because of changes in the local fishing policy..

For the US Virgin Island site, we carried out a sinple linear regression on the log transformed nests to estimate the instaneous rate of population change. The estimate was 0.078 (S.E.= 0.0135).

For the Florida sites, error searate error variance lognormal no 0.12579085 0.05233020 6 2.40 0.0530 lognormal yes 0.1215 0.0526 6 2.31 0.0603 Poisson no 0.1145 0.0423 6 2.71 0.0353 Poisson yes 0.1304 0.0569 6 2.29 0.0616 gamma no 0.1215 0.0526 6 2.31 0.0603 gamma yes 0.1167 0.0517 6 2.26 0.0647

For South America lognormal no -0.19018932 0.05958334 2 -3.19 0.0857 lognormal yes -0.16291771 0.04124533 2 -3.95 0.0585

References

Bijnens, L., Collette, L., Ivanov, A., Boes, G. H., and Sylvester, R. 1996. Optimal graphical display of the results of meta-analyses of individual patient data. 1996 Cochrane Colloquium.

Box, G. E. P., and Tiao, G. C. 1973. Bayesian Inference in Statistical Analysis. Wiley, New York.

Daniels, M. J. 1999. A prior for the variance in hierarchical models. The Canadian Journal of Statistics *27*(3): 567–578.

Galbraith, R. F. 1988. A note on graphical presentation of estimated odds ratios from several clinical trials. Statistics in Medicine **7**: 889–894.

Hedges, L. V., and Olkin, I. 1980. Vote-counting methods in research synthesis. Psychological Bulletin **88**: 359–369.

Hedges, L. V., and Olkin, I. 1985. Statistical Methods for Meta-analysis. Academic Press, San Diego.

Larivière, S., Jolicoeur, H., and Crête, M. 2000. Status and conservation of the gray wolf (*Canis lupus*) in wildlife reserves of Québec. Biological conservation **94**: 143–151.

Light, R. J., Singer, J. D., and Willett, J. B. 1994. The visual presentation and interpretation of meta-analyses. *In* The Handbook of Research Synthesis. *Edited by* H. Cooper and L. V. Hedges. Russell Sage Foundation, New York, pp. 439–453.

Olkin, I. 1999. Diagnostic statistical procedures in medical meta-analysis. Statistics in Medicine **18**: 2331–2341.

Searle, S. R., Casella, G., and McCulloch, C. E. 1992. Variance Components. John Wiley and Sons, New York, NY.

Smith, T. C., Spiegelhalter, D. J., and Thomas, A. 1995. Bayesian approaches to random-effects meta-analysis: a comparative study. **14**: 2685–2699.

Wolfinger, R., and O'Connell, M. 1993. Generalized linear mixed models: a pseudo-likelihood approach. J. Statist. Comput. Simul. **48**: 233–243.

8   Introduction

General Stuff

9   Example: Wolves in Québec

Over the last 100 years, the geographical range of gray wolves (*Canis lupus*) has been greatly reduced (Larivière, Jolicoeur, and Crête 2000). During the 1970's, the province of Québec established a network of wildlife reserves where the harvest of game species, including wolves, became controlled by a system of registered traplines and a quota on hunting licenses. Recently, Larivière, Jolicoeur, and Crête (2000) reported on wolf population trends during the last 15 years in 9 reserves located in southern Québec using data from questionnaires distributed to moose hunters and an equation linking the questionnaire data to radio-tracking data. In one of the reserves, Ashuap-mushuan, a different management scheme was used, and we omit the data from this reserve. Figure **??** displays the data from the remaining 8 reserves.

Note that the 1989 observation in the Mastigouche reserve seems unusual: it the lowest observation and deviates from the overall trend. Because it is not at either extreme, it might not be expected to be very influential on estimates of the intercept or slope in a linear regression, however it would be expected to inflate the variance estimate.

There is a prior reason for treating the 1997 observation in the Laurentides reserve with caution: the management scheme that year was changed as an experiment.

To estimate long-term trends in wolf densities in each reserve, Larivière, Jolicoeur, and Crête (2000) performed a simple linear regression of density versus year, and concluded that, over the last 15 years, in 7 of the 9 reserves, wolf populations have been relatively stable, with the remaining 2 reserves showing declines. This approach—determining the proportion of studies individually exhibiting statistical significance in the prescribed direction—is known as "vote counting" in the meta-analysis literature, and is inherently flawed (Hedges and Olkin 1985). Hedges and Olkin (1980) showed that as the number of studies becomes large, the proportion of studies yielding significant results is approximately the average power of the test.

10 Example: Wolves in Québec

Larivière, Jolicoeur, and Crête (2000) performed linear regressions on each of the data sets individually. For the ith data set, the model is

$$y_i = X_i\beta + \varepsilon_i, \tag{1}$$
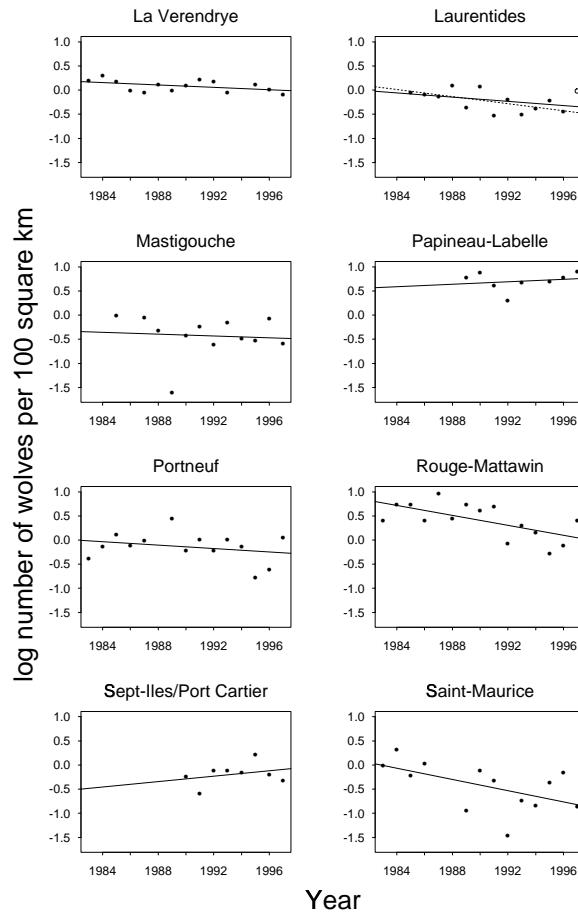
where $\varepsilon_i \sim N(0, \sigma_i^2 I)$.

Figure 2: Log population numbers of wolves in 8 reserves in southern Québec with fitted least squares regression lines. The 1997 observation in the Laurentides reserve is plotted as an open circle because the management scheme that year was changed. The dashed line shows the regression excluding that point.

Only three of the reserves give slopes significantly different from zero at the 95% confidence level—Laurentides, Rouge Mattawin, and Saint Maurice—and all three show declines. Two of the reserves show apparent (but non-significant) increases: Papineau-Labelle and Sept-Iles/Port Cartier.

Olkin (1999) encourages the meta-analyst to "Plot, plot, plot whenever and whatever you can." A standard display used in meta-analysis consists of a sequence of point estimates and confidence intervals from individual studies followed by a meta-analytic summary, typically a combined estimate with a confidence interval (Light et al. 1994; Galbraith 1988). This has also been called a forest plot (Bijnens et al. 1996). In subsequent chapters of this work, we will consider meta-analytic summaries, but for now, the focus is on individual estimates. Figure 3 shows a display of the individual estimates of the slopes of the regression lines in Figure 2.
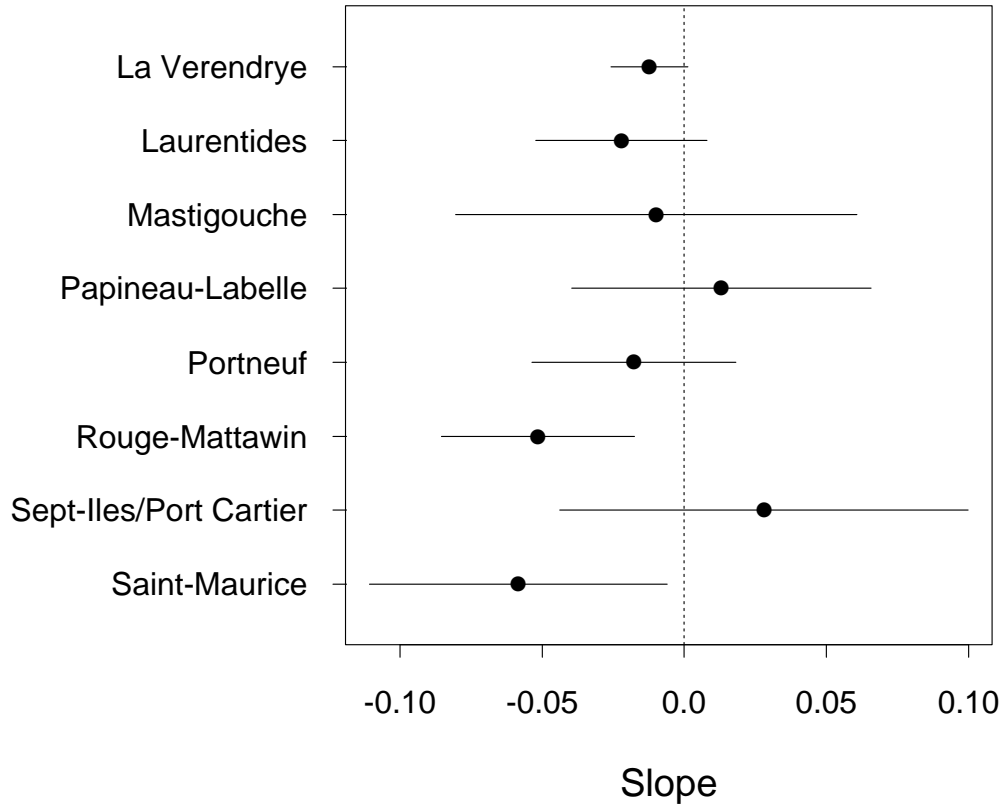
Figure 3: Point estimate and 95% confidence interval for the slope of the least squares regression of log numbers of wolves versus year for each reserve.

Figures 3 and 2 show completely independent reserve-specific estimates. In particular, the error variances, σ_i^2, are assumed to be unrelated. As Figure 2 shows, the error variances differ between reserves. For example, the error variance for the La Vérendrye reserve is clearly much smaller than that for the Saint-Maurice reserve. It may be that the true error variances differ between studies solely because of differences in the effective sample size underlying each observation from the various reserves. Approximate effective sample sizes, n_i^*, for observations from each reserve were obtained by ... and are shown in Table 1.

Reserve	Effective sample size, n_i^*
La Vérendrye	378
Laurentides	355
Mastigouche	168
Papineau-Labelle	198
Portneuf	100
Rouge-Mattawin	130
Sept-Îles/Port-Cartier	48
Saint-Maurice	66

Table 1: Effective sample sizes for observations from each reserve.

If we assume that the error variance for each reserve is given by σ^2/n_i^*, then we can estimate a single variance parameter, σ^2, thereby linking the 8 regressions. The resulting slope estimates are no longer "individual estimates" in the sense used above, but the assumption is relatively weak, and is unlikely to distort our conclusions seriously.

10.1   Example: Wolf data

As noted earlier there is at least one suspicious observation: the 1989 observation at Mastigouche. A robust regression procedure applied to these data can be helpful in assessing the sensitivity of the individual regression estimates, and in identifying unusual observations (Figure 4).
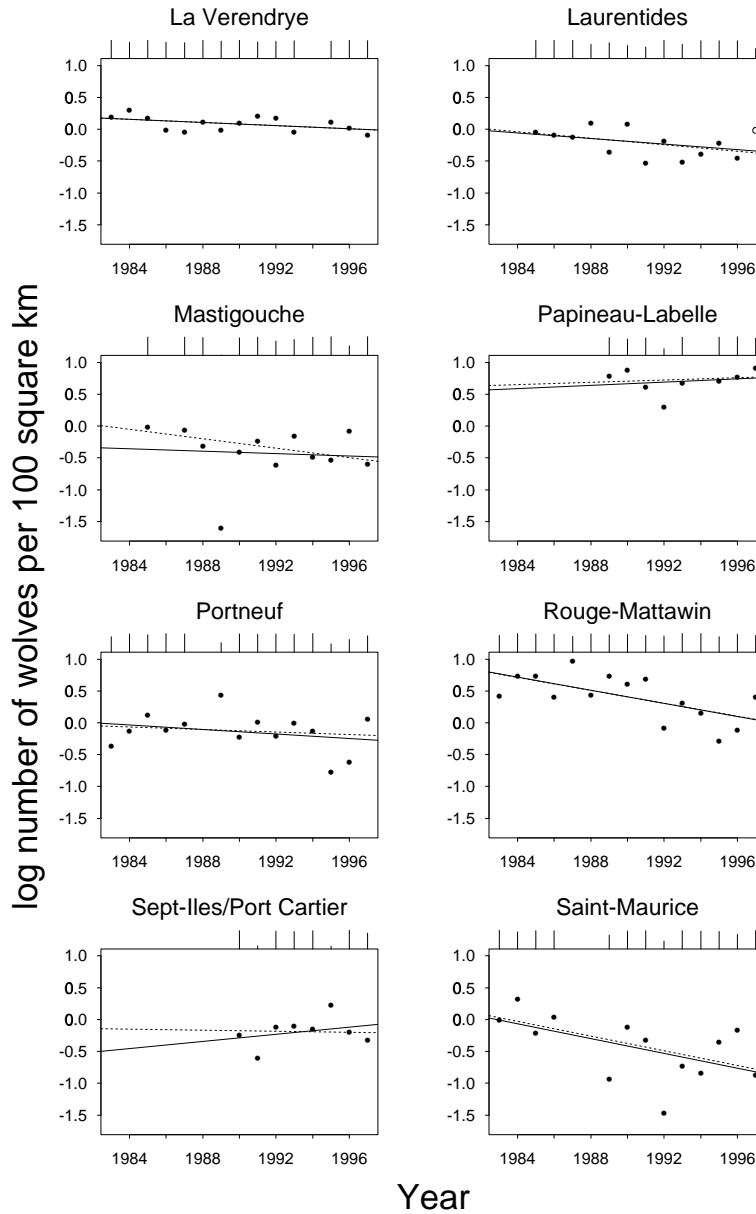
Figure 4: Log population numbers of wolves in 8 reserves in southern Québec with fitted least squares regression lines (solid) and robust regression lines (dotted). Along the top of each panel short vertical lines indicate the weights for each observation from the robust regression. When an observation is completely downweighted, no vertical line is visible. The 1997 observation in the Laurentides reserve is plotted as an open circle because the management scheme that year was changed.

Note that the 1989 observation at Mastigouche has been entirely eliminated from the analysis. Together with the downweighting of several other points for this reserve, the result is a substantial change in the estimated slope. Several observations at other reserves are also strongly down-weighted, e.g. the 1997 observation at Laurentides (which we have a prior reason to reject) and

12

the 1992 observation at Papineau-Labelle. Two points, 1991 and 1995, are downweighted in the Sept-Iles/Port Cartier data, resulting in a substantial change in slope. For such small data sets, it is not clear whether to accept the results from the robust regression. Nevertheless, it helps to highlight unusual observations and sensitivity in our estimates.

11 Example: linear mixed effects analysis of the wolf data

The reserves vary in size and density of wildlife, and these properties can be thought of as intrinsic to each reserve, and unrelated to any other reserves . Since the intercept for each reserve determines the initial density, we treat the intercepts as reserve-specific fixed effects. The rate of change of the population in each reserve, however, may reflect larger-scale phenomena such as climate or management policies. Therefore we model the slope for each reserve as a random effect.
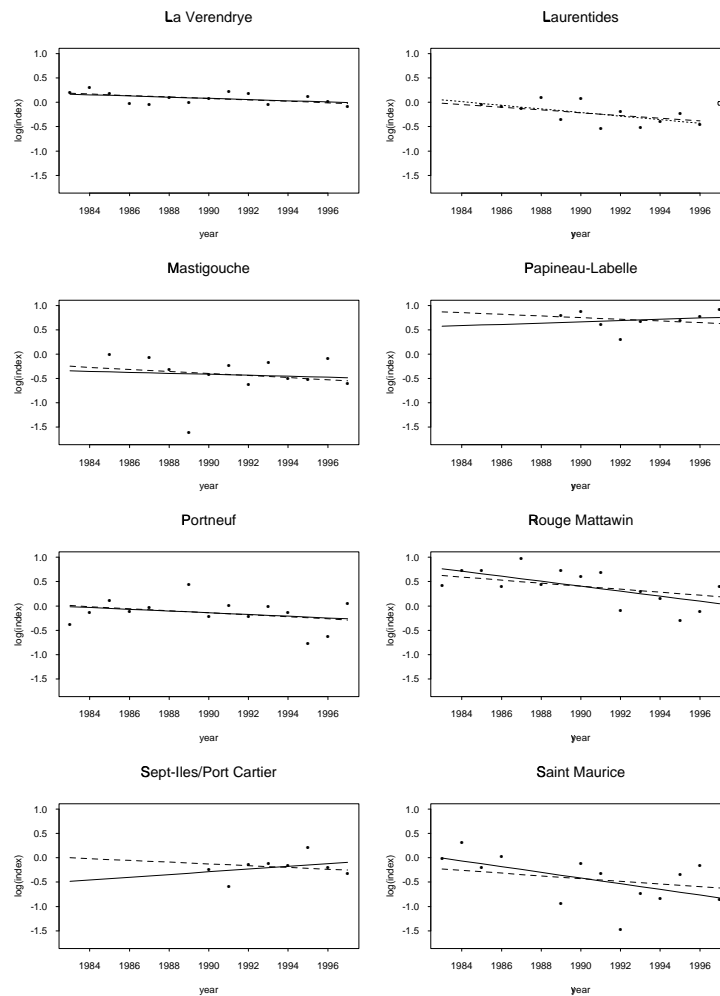

Figure 5: Log population numbers of wolves in 8 reserves in southern Québec with fitted least squares regression lines (solid) and BLUPs (dashed).

Reserve-specific intercepts
random slopes
reserve-specific error variances

SAS model

```
model logn = site year/solution;
random year /subject=site solution type=simple;
repeated /subject=site group=site;
```

SAS Output

```
            Covariance Parameter Estimates (REML)


     Cov Parm  Subject  Group           Estimate

     YEAR      SITE                      0.00013567
     DIAG      SITE     SITE LaVere      0.01237370
     DIAG      SITE     SITE Lauren      0.03215638
     DIAG      SITE     SITE Mastin      0.18672707
     DIAG      SITE     SITE PapLab      0.04536965
     DIAG      SITE     SITE Portne      0.08787437
     DIAG      SITE     SITE RouMat      0.08874428
     DIAG      SITE     SITE SeIPoC      0.06213305
     DIAG      SITE     SITE StMaur      0.19734598
```

The σ_i^2's clearly vary among reserves. In fact we do have proxies n_i^* for the sample sizes associated with the observations from the different reserves, shown in the table below

Reserve	LaVere	Lauren	Mastin	PapLab	Portne	RouMat	SeIPoC	StMaur
n_i^*	378	355	168	198	100	130	48	66

For example, each observation from LaVere is treated as being the average of 378 measurements. Writing $\sigma_i^2 = \sigma^2/n_i^*$ we now have a single error variance component to estimate.

Reserve-specific intercepts
random slopes
single error variance with reserve-specific weights

SAS model

```
model logn = site year/solution;
random year /subject=site solution type=simple;
weight sampsize;
```

SAS Output

```
        Covariance Parameter Estimates (REML)

 Cov Parm   Subject      Estimate

 YEAR       SITE        0.00000000
 Residual               12.55453776
```

The mean slope is estimated to be -0.026 with a standard error of 0.009. The variance of the slope is estimated to be zero. Using those point estimates suggests that none of the slopes could in fact be positive. However, this is a misleading conclusion, since it ignores the uncertainty in the estimates. Figure 6 is a graphical display of the range of plausible values of the mean and the standard deviation of the slope, with the associated probabilities of a positive slope. The "mushroom"-shaped contour is the joint 95% likelihood-based confidence interval for the mean and standard deviation of the slope. In the upper right portion of the mushroom, the mean is relatively large and the standard deviation is relatively large. The probability of a positive slope is then up to about 0.4.
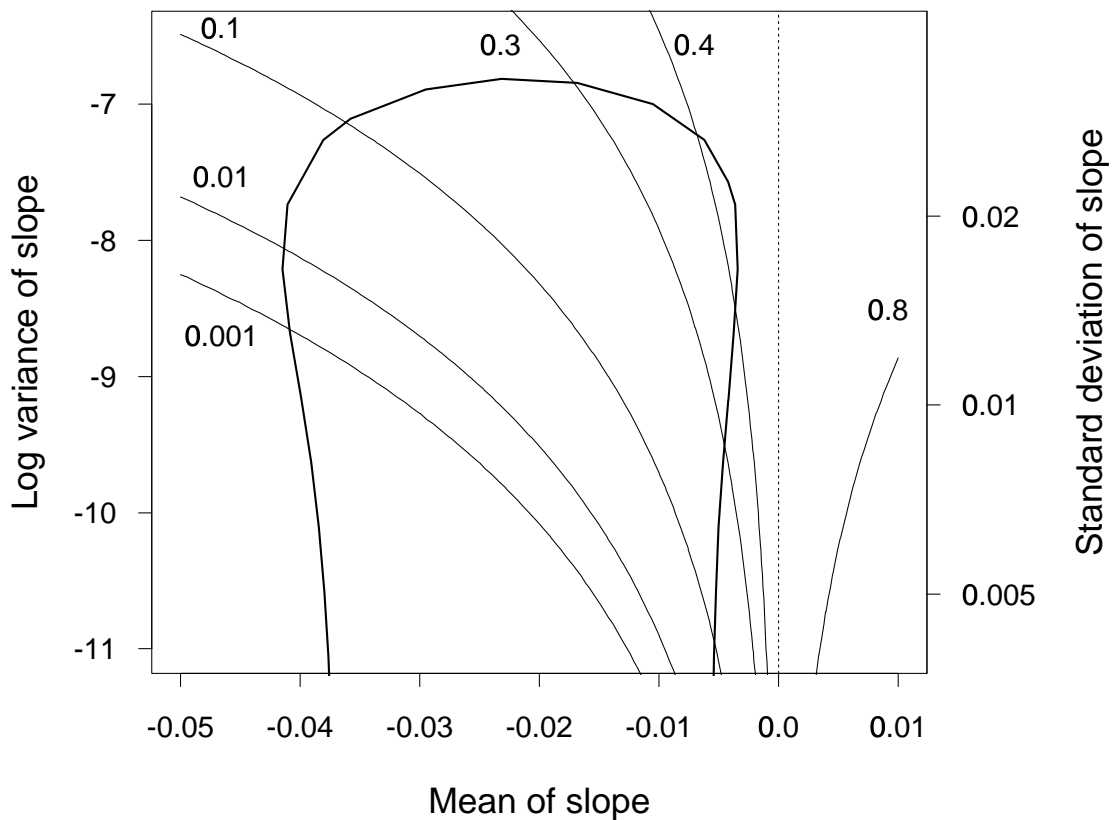

Figure 6: Approximate joint 95% confidence interval for $\log \sigma_1^2$ and $\log \beta_1$ (heavy line) with superimposed contours of constant probability (lighter lines), showing the probability of a positive slope. The dotted vertical line indicates a mean slope of zero.

A hierarchical Bayesian approach to this problem would make use of prior distributions for the mean and variance of the slope. But the choice of prior distributions for variances is quite difficult and in small samples "noninformative" priors may in fact be informative (Daniels 1999). Note first that the boundary value $\sigma^2 = 0$ is supported by non-negligible likelihood since it is possible that there is no between-study variance (cite DuMouchel). The standard noninformative prior is

16

$p(\sigma^2) \propto 1/\sigma^2$ (Box and Tiao 1973), however this has an asymptote at zero, which can lead to an improper posterior. Smith, Spiegelhalter, and Thomas (1995) argue that a reasonable approach is to search for a proper prior.

Profile log likelihood for SD of slope
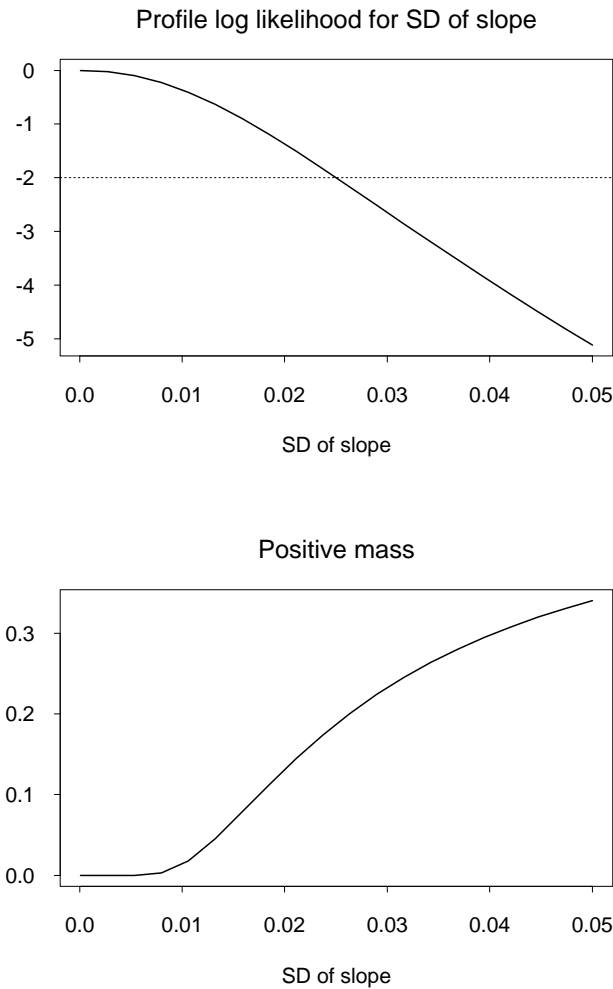
Positive mass

Figure 7: Profile log likelihood and ...

The general problem is to assess what is happening in a wide geographical area from limited, sampling in certain areas. This is similar to the problem in meta-analysis of synthesizing data on several experiments. The approach taken by NEF is to look at the proportion of study areas that show statistically significant results; this is known as "naive vote counting" in the meta-analysis literature (Hedges and Olkin 19??). There are grave disadvantages to this approach and can lead to very misleading results.

Surveys.

The wolf density was estimated by a combing information on the proportion of trips by moose hunters that observed tracts and scat was observed, and the number of days hunters heard wolf howls. The indices of abundance was transformed into density using the methods of ???. The

proportion data was arcsine transformed. Although the data is converted into numbers of wolves per 100 sq. kilometers, for our purpose, we only need to use the data as an index.

The province carried out 9 simple linear regressions of wolf density versus time, and concluded that that in general wolf populations were stable, because only 2 of the 9 reserves were found to have significant negative slopes.

Here we reconsider the question using a meta-analytic approach, to ask if the conclusions are reasonable. If the reserves are subject to the same management regime, and are subject to similar, large-scale environmental variability, then it is reasonable that to consider the hypothesis that they represent different "realizations" from a common distribution. Of the 9 reserves, 8 have a similar management regime and trapped primarily by non-natives, while one natives hold exclusive rights to trapping wolves. This last reserve, Ashaptashut, will be considered separately. It also is the only reserve that does not have data after 1993, the period when new, more efficient trapping methods were introduced (see next section). The other management change that occured was that in the last year for which we had data for Laurentide reserve most trappers voluntarily refrained from capturing wolves to assist in the study described in the next section. We thus, removed this point (1997) from the analysis.

We ask if we consider the data as a whole, including the errors in the estimates of abundance, weather we can conclude any thing about population change of the wolf populations.

We begin with the simplest of models. Let x_{it} be the estimate of abundance in reserve i in year t. We consider the initial model that each population is changing at a constant rate of the period of time of the censuses, we have

$$x_{it} = x_{i0}exp(r_i t + \varepsilon_{it})$$

where r_i is the rate of population change, x_{i0} is the initial population size of the ith population, and ε_{it} is the error is the estimate of abundance and deviations from the assumed model (sometimes called "process error"). It is unlikely that all the reserves will have exactly the same rate of change, so we will investigate the a simple mixed effect model where we assume that $r_i \sim N(\mu, \sigma_r)$. We estimated the parameters of this model under two different assumptions for ε_{it}. First, we assumed that the variance of ε_{it} is proportional to the inverse of the sample size, which we approximately know. That is , we assume that $\varepsilon_{it} \sim \sigma/n_i$, where σ is estimated from the data. An alternative approach is to assume that the variances are primarily due to factors that are not proportional to sample size, and estimate a separate variance, σ_i for each reserves.

Results

When fit individually, there are 2 positive slopes estimated, and the rest negative, with statistically significant slopes. These results are similar to the MEF's analysis without the log transform. Both versions of the mixed model analysis show a statistically significant decline in the mean (μ) of around 0.02 (roughly 2% a year), and estimate very little variability among reserves in the rate of population change. This suggests that the 8 reserves are behaving a similar fashion, and are probably representative of much of southern Quebec where trapping of wolves occur.

Table

We also tested if autocorrelation in the residuals was important. We thus fit the same model as

18

before but we assumed that ε_{it} was a first-order autoregressive process for each reserve.

References

Bijnens, L., Collette, L., Ivanov, A., Boes, G. H., and Sylvester, R. 1996. Optimal graphical display of the results of meta-analyses of individual patient data. 1996 Cochrane Colloquium.

Box, G. E. P., and Tiao, G. C. 1973. Bayesian Inference in Statistical Analysis. Wiley, New York.

Daniels, M. J. 1999. A prior for the variance in hierarchical models. The Canadian Journal of Statistics **27**(3): 567–578.

Galbraith, R. F. 1988. A note on graphical presentation of estimated odds ratios from several clinical trials. Statistics in Medicine **7**: 889–894.

Hedges, L. V., and Olkin, I. 1980. Vote-counting methods in research synthesis. Psychological Bulletin **88**: 359–369.

Hedges, L. V., and Olkin, I. 1985. Statistical Methods for Meta-analysis. Academic Press, San Diego.

Larivière, S., Jolicoeur, H., and Crête, M. 2000. Status and conservation of the gray wolf (*Canis lupus*) in wildlife reserves of Québec. Biological conservation **94**: 143–151.

Light, R. J., Singer, J. D., and Willett, J. B. 1994. The visual presentation and interpretation of meta-analyses. *In* The Handbook of Research Synthesis. *Edited by* H. Cooper and L. V. Hedges. Russell Sage Foundation, New York, pp. 439–453.

Olkin, I. 1999. Diagnostic statistical procedures in medical meta-analysis. Statistics in Medicine **18**: 2331–2341.

Searle, S. R., Casella, G., and McCulloch, C. E. 1992. Variance Components. John Wiley and Sons, New York, NY.

Smith, T. C., Spiegelhalter, D. J., and Thomas, A. 1995. Bayesian approaches to random-effects meta-analysis: a comparative study. **14**: 2685–2699.

Wolfinger, R., and O'Connell, M. 1993. Generalized linear mixed models: a pseudo-likelihood approach. J. Statist. Comput. Simul. **48**: 233–243.