

Modelling paired release–recovery data in the presence of survival and capture heterogeneity with application to marked juvenile salmon

Ken B Newman

Centre for Research into Ecological and Environmental Modelling, University of St Andrews, St Andrews, Scotland

Abstract: Products of multinomial models have been the standard approach to analysing animal release–recovery data. Two alternatives, a pseudo-likelihood model and a Bayesian nonlinear hierarchical model, are developed. Both approaches can to some degree account for heterogeneity in survival and capture probabilities over and above that accounted for by covariates. The pseudo-likelihood approach allows for recovery period specific overdispersion. The hierarchical approach treats survival and capture rates as a sum of fixed and random effects. The standard and alternative approaches were applied to a set of paired release–recovery salmon data. Marked juvenile chinook salmon (*Oncorhynchus tshawytscha*) were released, with some recovered in freshwater as juveniles and others in marine waters as adults. Interest centered on modelling freshwater survival rates as a function of biological and hydrological covariates. Under the product multinomial formulation, most covariates were statistically significant. In contrast, under the pseudo-likelihood and hierarchical formulations, the standard errors for the coefficients were considerably larger, with pseudo-likelihood standard errors five to eight times larger, and fewer coefficients were statistically significant. Covariates, significant under all formulations, with important management implications included water temperature, water flow and amount of water exported for human use. The hierarchical model was considerably more stable with regard to estimated coefficients of training subsets used in a cross-validation.

Key words: band-recovery models; hierarchical; Markov chain Monte Carlo; mixed effects; pseudo-likelihood

Data and software link available from: <http://stat.uibk.ac.at/SMIJ>

Received May 2001; revised June 2002, January 2003; accepted April 2003

1 Introduction

Release–recovery data are generated by marking animals, releasing them, and later recovering them. In contrast to capture–recapture data, the recovered animals are not rereleased; in many cases the recovered animals are dead. Band–recovery data where banded birds are later recovered by hunters are an example of release–recovery data, as are marked and tagged salmon recovered by fisheries. The classic approaches to modelling release–recovery and capture–recapture data date back to the 1950s and

Address for correspondence: Ken B Newman, Senior Lecturer, CREEM, University of St Andrews, The Observatory, Buchanan Gardens, St Andrews KY16 9LZ, Scotland. E-mail: ken@mcs.st-and.ac.uk

early 1960s (Darroch, 1959; Cormack, 1964; Jolly, 1965; Seber, 1965) and are based upon products of multinomial models. In the case of release–recovery data with a sequence of possible recovery points, the marginal distribution for recoveries at a given time point is binomial with the recovery probability being the product of a sequence of conditional survival and conditional capture probabilities. For example, let R be the number of animals released at the beginning of the study, S_i be the probability of surviving to time point t_i (given it was alive at t_{i-1}), and p_i be the probability that an animal alive at t_i is captured then. The number of recoveries at time point t_3 , say, is $\text{Binomial}(R, S_1(1 - p_1)S_2(1 - p_2)S_3p_3)$.

Extensions to the multinomial models have included the modelling of survival and capture probabilities by covariates and the inclusion of an overdispersion parameter to account for extra-multinomial variance (Lebreton *et al.*, 1992). Departures from the multinomial distributions include Poisson distributions, Poisson generalized linear models, and overdispersed Poisson or quasi-likelihood models (Cormack, 1993).

Presented in this paper are two further extensions to the classic multinomial formulations. One, labelled the pseudo-likelihood approach, allows for different overdispersion parameters at different recovery points. An earlier analysis of a set of unpaired release–recovery salmon data (Newman and Rice, 2002) used extended quasi-likelihood (Nelder and Pregibon, 1987) to fit two overdispersion parameters at two recovery points. This paper is a slight variation over this earlier work in that the data are a set of paired release–recovery data and pseudo-likelihood (Carroll and Ruppert, 1988) is used for estimating two overdispersion parameters.

The second extension to the classic formulation is a nonlinear hierarchical model where survival and recovery probabilities for a given release of marked animals are a function of fixed and random effects. Heterogeneity in survival and recovery probabilities that cannot be explained by covariates is then partially accounted for by the random effects.

This work was motivated by an application to release–recovery chinook salmon data and the next section describes the motivation of the application, the data and reasons for overdispersion. The following section describes the three modelling approaches and estimation procedures. The results are then compared and the paper ends with a discussion that includes a sensitivity analysis, more complex hierarchical models, a comparison with the unpaired release–recovery analysis (Newman and Rice, 2002), and brief discussion on choosing between the three alternative modelling approaches.

2 Application background

The Sacramento River is located in northern and central California and provides water for human consumption and agricultural use to over 20 million people. The river is also home to chinook salmon (*Oncorhynchus tshawytscha*), a species of salmon that had returns of a million or more in the early 1900s (Healey, 1991). In the last 30 years there has been a drastic reduction in the number of naturally spawning salmon due to loss of habitat, environmental degradation and overfishing. The loss of natural fish has been somewhat mitigated by hatchery-produced salmon, but considerable concern remains over the viability of naturally spawning stocks.

To identify which factors influence the survival of juvenile chinook salmon as they outmigrate from freshwater to marine waters, the US Fish and Wildlife Service (USFWS) has been conducting release-recovery studies for over 20 years. Hatchery-reared juvenile chinook salmon are externally marked by removing the adipose fin and a micro-tag (a coded-wire-tag) is injected into the snout. The tags are release or batch specific, not individual fish specific, and the fish must be sacrificed to read the tag and identify the release the fish came from. The fish are released at multiple locations in the Sacramento River, particularly in the lower portions of the river, during the months of April and May. A midwater trawl located downstream of the release sites, in the tidal zone just east of San Francisco Bay, recovers the fish within two to three weeks after release.

For some of the study years additional releases of marked and tagged fish were made just downstream of the trawl and these releases are viewed as being paired with some of the upstream releases. Fish from both the upstream and downstream release locations are later caught as two- to five-year old fish in the Pacific Ocean by commercial and recreational fisheries. At landing ports throughout the fishing season, samples are taken of the catches, additional recoveries of the marked and tagged fish are made, and estimates of the total number of marine fisheries' recoveries are made.

As mentioned previously, the river and marine recoveries from the upstream releases alone were analysed by Newman and Rice (2002), here labelled the unpaired releases analysis. Let S be the probability of surviving from point of release to the trawl and p be the conditional probability of capture given survival. Sp was modeled as a function of several biological and hydrological covariates. Using the upstream releases alone, S and p are not separately estimable. However, by assuming that the capture rate p was a product of known trawl fishing effort f and some unknown, but constant catchability coefficient, q , the ratio S_1/S_2 for any two releases could be estimated. Letting $\widehat{S}_i p_i$ be the model-based estimate of Sp and f_i be the corresponding trawl fishing effort for release group i , $i = 1, 2$,

$$\frac{\widehat{S}_1}{\widehat{S}_2} = \frac{\widehat{S}_1 p_1 / f_1}{\widehat{S}_2 p_2 / f_2} \quad (2.1)$$

In contrast with Newman and Rice (2002), this paper presents an analysis of recoveries from the paired upstream and downstream releases. By assuming identical marine survival and capture rates for paired upstream and downstream releases, S and p are separately estimable. The release and recovery data along with covariates used to model S are described next, followed by a section listing reasons for overdispersion with these data.

2.1 Data

With assistance from USFWS personnel, 61 upstream releases were paired with 19 downstream releases made between 1979 and 1995. The term 'pairings' is not strictly correct in that several upstream releases were sometimes matched with a single common downstream release. The term release set will sometimes be used alternately with release pair. Most of the upstream releases were among the larger set of 101 upstream releases analysed by Newman and Rice (2002). The upstream release groups were released in

one of three locations over the period of study, near the city of Sacramento (approximately 50 miles upstream of the trawl; $n = 22$), near Courtland (38 miles upstream; $n = 18$), and near Ryde (30 miles upstream; $n = 21$). The downstream groups were released at one of two locations below the trawl, Port Chicago and Benicia.

Table 1 summarizes the release and recovery information. The number of fish released from an upstream location is denoted R_u , while the number released downstream is R_d . The number of trawl recoveries from an upstream release is y_{ut} , the number of ocean recoveries from an upstream release is y_{uo} , and the number of ocean recoveries from a downstream release is y_{do} . As mentioned earlier the number of ocean recoveries is estimated by what is approximately a temporally-spatially stratified sample of marine catch. These estimates, denoted \hat{y}_{uo} and \hat{y}_{do} ; can be approximately written as

$$\hat{y}_{uo} \approx \sum_a \sum_t \sum_p e_{atp} y_{uatp}, \quad \hat{y}_{do} \approx \sum_a \sum_t \sum_p e_{atp} y_{datp}$$

where y_{uatp} and y_{datp} are the number of recoveries, from an upstream release and a downstream release, of marked and tagged fish in stratum atp 's sample, where a denotes age (which ranges from 2 to 5 years), t denotes time period within a fishing season, and p denotes landing area (usually a port). e_{atp} is the inverse of the sampling fraction for a given stratum, also known as the expansion factor, and is on average 4 to 5, that is, 20 to 25% of the landed catches are sampled. The trawl recovery rates, r_{ut} , and estimated ocean recovery rates, \hat{r}_{uo} and \hat{r}_{do} are defined as the ratio of recoveries, or estimated recoveries, to number released, that is,

$$r_{ut} = \frac{y_{ut}}{R_u}, \quad \hat{r}_{uo} = \frac{\hat{y}_{uo}}{R_u}, \quad \hat{r}_{do} = \frac{\hat{y}_{do}}{R_d}.$$

For upstream releases the median \hat{r}_{uo} was an order of magnitude greater than the median r_{ut} . For downstream releases the median \hat{r}_{do} was about twice \hat{r}_{uo} . For a paired release, assuming common ocean survival, harvest, and sampling rates, r_{uo} should be less than r_{do} due to in-river mortality and removals by the trawl.

Table 1 Summary of release and recovery information for 61 upstream releases and 19 downstream releases. R_u and R_d are the number of fish released either upstream or downstream of the trawl. The number of upstream fish recovered by the trawl are denoted y_{ut} , while \hat{y}_{uo} and \hat{y}_{do} are the estimated number of recoveries in the ocean fisheries of upstream and downstream releases, respectively. The fraction of recoveries to releases are denoted r_{ut} , \hat{r}_{uo} , and \hat{r}_{do} .

	Min	1st Qu.	Median	3rd Qu.	Max
R_u	10 887	50 601	51 819	57 561	160 151
y_{ut}	2	14	35	67	145
\hat{y}_{uo}	10	98	360	562	1979
r_{ut}	0.00005	0.00032	0.00059	0.00102	0.00272
\hat{r}_{uo}	0.00020	0.00397	0.00536	0.01050	0.02485
R_d	42 000	48 069	54 055	71 332	110 122
\hat{y}_{do}	129	359	782	1136	3338
\hat{r}_{do}	0.00297	0.00739	0.01074	0.02070	0.03241

Table 2 Summary of covariates used in modelling of *S*

Covariate	Description	\bar{x}	<i>s</i>
Size	Average length in mm	80.92	6.11
Log flow	Log transformed median river flow in cfs during the outmigration period	9.53	0.45
Salinity	Water salinity as measured by resistance, μ mho/cm	5219.79	3756.16
Release temperature	River water temperature ($^{\circ}$ F) at release	65.71	4.75
Hatchery temperature	Water temperature ($^{\circ}$ F) in hatchery on day of release	54.55	3.04
Tide	A measure of the magnitude of the change in low-low and high-low tides and whether the delta was filling or draining	1.59	0.70
Exports	Median volume of water in cfs diverted during the outmigration period	4888.23	2141.72
Gate	Indicator for position of the cross-channel gate located just below Courtland; 1 if open and 0 if closed	0.61	0.49
Turbidity	Turbidity of water (formazine turbidity units)	8.18	3.70

\bar{x} and *s* are the mean and sample standard deviation.

The covariates used to model *S* are summarized in Table 2. The covariate values for each of the three upstream release sites, individually, were quite similar in terms of means and standard deviations. The gate variable is an indicator for the position of a diversion gate located just downstream of Courtland. When the gate is open (indicator = 1), fish moving downstream are more likely to get diverted into a sprawling delta where large water export pumps are located. Indicator variables for release at Sacramento or Courtland (labelled *Sac* and *Court*) were included in the modelling of *S*. This allowed for a release site effect that was to some degree a function of distance upstream. Consistent with Newman and Rice (2002), gate position and export level were assumed to only affect releases made above the diversion gate (Sacramento and Courtland alone). Interaction terms, crossing the Sacramento and Courtland indicator variables with exports and gate position, were used to reflect that assumption; that is, the gate and export interaction variable values were set at 0 for Ryde releases.

Correlations between covariates were slight (less than 0.5) with two exceptions. Flow and salinity are inversely related ($r = -0.74$), but not in a strictly linear fashion; as outflow increases, the influx of seawater lessens. Hatchery and release temperatures are positively correlated ($r = 0.67$), because the water source for the hatchery is river water.

The covariates are a subset of those used by Newman and Rice (2002) to model *Sp*, with the exception that here the export measure is total volume exported, while Newman and Rice used the ratio of export volume to flow volume; subsequent differences are discussed later.

2.2 Reasons for overdispersion

The simplest way to model recoveries is to assume that all the fish in a given release are independent and have the same probabilities of recovery (by the trawl or by the ocean

fishery). In other words for a given upstream release, trawl and ocean recoveries are trinomial, and for a given downstream release, ocean recoveries are binomial. Such an approach, labelled the trinomial/binomial product (TBP) model, is the basis for one of the approaches taken and is discussed in the next section. The fact that ocean recoveries are estimated rather than observed make the TBP formulation questionable, however, and make overdispersion likely. Letting π be the probability that a downstream release is later caught by the ocean fishery, the variation in \hat{y}_{do} is greater than for a Binomial(R_d, π) random variable due to estimation error. Theoretical arguments and empirical evidence for overdispersion led to the two alternative modelling approaches.

There are several reasons for possible overdispersion (Collett, 1991). Important covariates, either unknown or unmeasured, may have been left out of the model for the recovery probability. To lessen the chance of this, a relatively conservative approach was taken to the modelling of S , p and π and is described later. Overdispersion is likely due to correlation between individuals, caused by fish schooling or clustering, in particular while moving downstream. Additionally, heterogeneity in survival probabilities is likely due to variation in individual fish size at least. Relatedly, group-level covariate values, such as average fish length, are used in the modelling of survival, thus covariate values have measurement error, which induces correlation between individuals and subsequent overdispersion (Prentice, 1986). Heterogeneity in survival probabilities is also likely due to fish from the same release taking different routes downstream and having different travel times to the trawl. The trawl operates only during portions of the day and can only sweep a portion of the width and depth of the river when it is operating. Variation in travel times and position in the river then translates into heterogeneity in capture probabilities, too.

There was empirical evidence for overdispersion. Several upstream release groups were identified as replicates, groups nearly identical with the exception of having different tag codes. Based on a χ^2 goodness of fit test assuming a multinomial model, the variation in river and (estimated) ocean recoveries for some of the replicate sets, but not all, was greater than expected. For example, assuming a trinomial model for the downstream and ocean recoveries from four replicate releases from Courtland in 1985 yielded $\chi^2 = 25.8$, 6 df, P -value = 0.0002.

3 Methods

3.1 Tri/binomial product (TBP) model for recoveries

The tri/binomial product model is a particular case of a band-recovery or release-dead recovery model (Brownie *et al.*, 1985) based on products of multinomial distributions. Each upstream release has one of three possible fates, recovery by the trawl, recovery in the ocean fisheries and anything else. The fate for any individual fish is assumed independent of the fate any other fish. For a paired release, or release set, it is assumed that the ocean recovery probability, π , is the same for all releases. For a given release

pair, the joint distribution of (y_{ur}, y_{uo}, y_{do}) is a product of trinomial and binomial distributions:

$$\Pr(y_{ur}, y_{uo}, y_{do}) = \binom{R_u}{y_{ur} \quad y_{uo}} (Sp)^{y_{ur}} (S(1-p)\pi)^{y_{uo}} (1-Sp - S(1-p)\pi)^{R_u - y_{ur} - y_{uo}} \\ \times \binom{R_d}{y_{do}} \pi^{y_{do}} (1-\pi)^{R_d - y_{do}} \quad (3.1)$$

To address the management questions about the effect of biological and hydrological variables on survival probability, S was modelled as a function of the covariates (Table 2). Capture probabilities were handled two different ways: (a) 'release specific' ps (no covariates) and (b) covariate-based ps . Maximum likelihood estimates of the covariates for S , of p or its covariates, and π were calculated for the TBP model with estimates of y_{uo} and y_{do} substituted in equation (3.1) for the actual, but unknown values. The models for S and p were the same under the TBP, pseudo-likelihood and hierarchical formulations, and are described next.

3.2 Modelling S

The logit of S was modeled as a function of covariates (Lebreton *et al.*, 1992); that is,

$$\ln S = \frac{\exp(x'\beta)}{1 + \exp(x'\beta)} \quad (3.2)$$

where x is a column vector of covariates and

$$x'\beta = \beta_0 + \beta_1 Sac + \beta_2 Court + \beta_3 Size + \beta_4 Log Flow \\ + \beta_5 Salinity + \beta_6 Release Temp + \beta_7 Hatchery Temp + \beta_8 Tide \\ + \beta_9 (Sac|Court \times Exports) + \beta_{10} (Sac|Court \times Gate) + \beta_{11} Turbidity \quad (3.3)$$

Again, Sac and $Court$ are indicators for releases from Sacramento and Courtland. The covariates $(Sac|Court \times Exports)$ and $Sac|Court \times Gate$ are the exports and gate position indicator just for releases from Sacramento or Courtland (the values are zero for releases from Ryde). To facilitate comparisons between covariates in terms of the magnitudes of coefficients, and to lessen numerical errors, the nonindicator variables were standardized.

3.3 Modelling p

The river capture probabilities, ps , and the ocean recovery rates, π s, can be modeled as functions of covariates. There are several questions of interest one could answer by doing so; for example, 'How related is p to trawl effort?', 'Is p affected by flow or turbidity?' Only two models for p were examined (and none for π), however, because the primary focus was on factors affecting survival. Incorrect modelling of the ps or the

π s results in biased estimates of survival. Conversely, the cost of not modelling the p s or π s is less precision in the estimates of the survival model coefficients; using release set specific p s and π s adds 80 parameters.

One model for p used trawl effort as a covariate. The other model used an indicator for 1988 releases, the year in which sampling effort was approximately double that of other years:

$$p = \frac{\exp(\gamma_0 + \gamma_1 I_{1988})}{1 + \exp(\gamma_0 + \gamma_1 I_{1988})} \quad (3.4)$$

The indicator variable was selected over the effort measure because the former yielded a larger likelihood at the maximum likelihood value from the TBP model and it provided a slightly better fit as measured by a χ^2 goodness of fit statistic. Only results for equation (3.4) are reported.

3.4 Pseudo-likelihood (PL) model for recoveries

The pseudo-likelihood approach of Carroll and Rupert (1988) was used to account for overdispersion in both the trawl and estimated ocean recoveries. Different overdispersion parameters were used for each recovery type. The expected number of recoveries were based on the TBP formulation.

$$\begin{aligned} E[y_{ut}] &= R_u S p \\ E[\hat{y}_{uo}] &= R_u S (1 - p) \pi \\ E[\hat{y}_{do}] &= R_d \pi \end{aligned}$$

The variances were multiples of the TBP variances with three different dispersion parameters used for each recovery category.

$$\begin{aligned} \text{Var}[y_{ut}] &= \phi_{ut} R_u S p (1 - S p) \\ \text{Var}[\hat{y}_{uo}] &= \phi_{uo} R_u S (1 - p) \pi (1 - S (1 - p) \pi) \\ \text{Var}[\hat{y}_{do}] &= \phi_{do} R_d \pi (1 - \pi) \end{aligned}$$

In the absence of replicates among the downstream releases, overdispersion for y_{do} cannot be estimated and ϕ_{do} was fixed at 1.0. Similarly, when release-specific values p are used, ϕ_{ut} is not estimable and was fixed at 1.0. Even if replicates were available for upstream and downstream releases and all the ϕ s were estimable, the assumption of constant values for each release-recovery combination is at best a coarse means of dealing with the overdispersion. For example, if the survival and capture probabilities are viewed as random variables (as in the hierarchical formulation discussed later), then the magnitude of the variance inflation is a function of the number of fish released. The range of release numbers is considerable (Table 1), thus between-release variation in overdispersion could be large.

The objective function to be maximized is

$$\begin{aligned}
 P = & -0.5 \left[\sum_{i=1}^{n_u} \left\{ \frac{(y_{uti} - \mu_{uti})^2}{\text{Var}[y_{uti}]} + \log(2\pi \text{Var}[y_{uti}]) \right\} \right] \\
 & -0.5 \left[\sum_{i=1}^{n_u} \left\{ \frac{(\hat{y}_{uoi} - \mu_{uoi})^2}{\text{Var}[\hat{y}_{uoi}]} + \log(2\pi \text{Var}[\hat{y}_{uoi}]) \right\} \right] \\
 & -0.5 \left[\sum_{i=1}^{n_d} \left\{ \frac{(\hat{y}_{doi} - \mu_{doi})^2}{\text{Var}[\hat{y}_{doi}]} + \log(2\pi \text{Var}[\hat{y}_{doi}]) \right\} \right]
 \end{aligned}$$

where n_u and n_d are the number of upstream and downstream releases; unsubscripted π is the constant 3.14159.

This formulation ignores the correlation between y_{ut} and \hat{y}_{uo} . However, given the relatively small magnitude of S , p and π , the effect of the correlation was practically ignorable. The median estimated correlation between y_{ut} and \hat{y}_{uo} was -0.002 .

3.5 Hierarchical model for recoveries

The pseudo-likelihood formulation can be viewed as an approximation to a hierarchical model where the parameter combinations Sp , $S(1-p)\pi$ and π are random variables arising from three hyperdistributions. Because S , p and π appear in more than one combination, it is awkward, at best, to arrive at meaningful hyperdistributions for each of the three combinations. It is more natural, and likely more accurate, to view the individual survival, capture and ocean recovery rates as arising from separate hyperdistributions and that was the approach taken here.

The first stage of the hierarchy is the distribution of recoveries (observed and estimated) for a single upstream and downstream pair, which conditional on S , p and π , is *assumed* TBP (see equation (3.1)).

$$\begin{aligned}
 y_{ut}, \hat{y}_{uo} & \sim \text{Trinomial}(R_u, Sp, S(1-p)\pi) \\
 \hat{y}_{do} & \sim \text{Binomial}(R_d, \pi)
 \end{aligned} \tag{3.5}$$

For the second level of the hierarchy, survival rates were modelled according to a logistic-normal distribution (Hinde and Demetrio, 1998), as were capture rates when modelled as a function of release year. When capture rates were release-specific, the prior distribution was $\text{Uniform}(0, 0.01)$; similarly the priors for ocean recovery rates were $\text{Uniform}(0, 0.08)$. The upper bound of 0.01 on the prior for p was based upon the trawl effort measure.

$$\log\left(\frac{S}{1-S}\right) \sim \text{Normal}(x'\beta, \sigma_S^2) \tag{3.6}$$

$$\log\left(\frac{p}{1-p}\right) \sim \text{Normal}(\gamma_0 + \gamma_1 I_{1988}, \sigma_p^2), \quad \text{or } p \sim \text{Uniform}(0, 0.01) \quad (3.7)$$

$$\pi \sim \text{Uniform}(0, 0.1) \quad (3.8)$$

At the top of the hierarchy, parameters of the logistic normal distributions were modelled as follows:

$$\beta_i \sim \text{Normal}\left(0, \frac{\pi^2}{3 \times 12}\right), \quad i = 0, 1, \dots, 11 \quad (3.9)$$

$$\sigma_S^2 \sim \text{Exponential}(0.001) \quad (3.10)$$

$$\gamma_0 \sim \text{Normal}(-5.60, \frac{0.987}{2}) \quad (3.11)$$

$$\gamma_1 \sim \text{Normal}(0, \frac{0.987}{2}) \quad (3.12)$$

$$\sigma_p^2 \sim \text{Exponential}(0.001) \quad (3.13)$$

The hyperparameters were chosen with uniform distributions for S and p in mind. In particular, assuming that the prior for S was $\text{Uniform}(0, 1)$, then $\log(S/(1-S))$ follows a logistic distribution with mean 0 and variance $\pi^2/3$. The covariates on the right-hand side of (3.6) were standardized to have an average of 0 and a standard deviation of 1, with the exception of indicator variables. With the β s defined as in (3.9) and σ_S^2 as in (3.10), the (unconditional) expected value and variance of the left-hand side of (3.6) are nearly 0 and $\pi^2/3$ (with some deviation due to values of nonstandardized covariates). Assuming a prior for the capture probabilities p of $U(0, 0.01)$, the priors for γ_0 , γ_1 and σ_p^2 were chosen by a similar procedure.

The hyperparameters for σ_S^2 and σ_p^2 , here 0.001, were selected in a nonstandard, but pragmatic manner. Three-fold cross-validation was conducted (discussed in detail later in the paper) and the average absolute errors were compared for different values of the hyperparameters, with the value 0.001 yielding small average errors. When first fitting the models the hyperparameters were chosen such that $E[\sigma_S^2] = \text{Var}[\beta]$ and $E[\sigma_p^2] = \text{Var}[\gamma]$. The resulting posterior distributions were such that the average posterior fitted values were very close to the observed values with the random effects relatively large. The predictive accuracy on the test sets of the cross-validation was relatively low, however, suggesting that the training data sets were being overfit.

3.6 Model fitting

For the TBP and the PL formulations, the objective functions, the log-likelihood and the pseudo-likelihood, were directly maximized using the automatic differentiation optimization program, AD Model Builder (Otter Research Ltd., Sidney, BC, Canada). AD Model Builder was also used to estimate the dispersion parameters of the PL model and to calculate covariance matrices.

Markov chain Monte Carlo (MCMC), in particular the Metropolis-Hastings algorithm, was used to generate samples from the posterior distributions for the

hyperparameters ($\beta_s, \gamma_s, \sigma_s, \pi_s$), the survival, trawl capture, and ocean recovery probabilities. Candidate values were generated in a block-like manner for the survival, capture and ocean recovery parameters, respectively. The `gibbsit` program of Raftery and Lewis (1996) was used to determine burn-in time and chain length. The proposal distributions were tuned such that a chain length of 40 000 was sufficient (with a burn-in of 2000 more than adequate). `gibbsit` was also used to determine the degree of thinning of MCMC output to allow use of standard nontime-series estimation of standard deviations for the hyperparameters' posterior distributions.

The survival model coefficients, the β_s in (3.6), were generated individually using a Metropolis proposal, $\beta_{\text{candidate}} \sim \text{Normal}(\beta_{\text{current}}, \sigma_{\text{tune}}^2)$. A candidate value for σ_s^2 was generated simultaneously with each individual β using a lognormal perturbation from the previous values, $\ln(\sigma_{s,\text{candidate}}^2) \sim \text{Normal}(\ln(\sigma_{s,\text{current}}^2), \sigma_{\text{tune}}^2)$. In both cases the variances of the normal distributions were tuned to yield reasonable acceptance and apparent mixing rates. After each 'new' value of β_i and σ_s was determined, candidate values for the individual release survival parameter, $S_j, j = 1, \dots, 61$, were generated on an individual basis using uniform proposal distributions centered around the previous value, with interval width used for tuning. A similar procedure was used for generating $\gamma_1, \gamma_0, \sigma_p$ and the trawl capture rates. For the ocean recovery probabilities, the π_s , release set-specific values were generated individually.

The algorithm is sketched below for the case of covariate-based p . At iteration t ,

- 1) For i in 0:11 {generate β_i and σ_s^2 ; for j in 1:61 {generate S_j }
- 2) For i in 0:1 {generate γ_i and σ_p^2 ; for j in 1:61 {generate p_j }
- 3) For j in 1:19 {generate π_j }

4 Results

4.1 Comparison of models

Estimates of the coefficients of the logistic model for S based on the TBP, PL and hierarchical formulations are shown along with standard errors in Table 3. The hierarchical model point estimates are the means of the posterior distributions and the standard errors are based on the thinned chain (chain length was 40 000, burn-in was 2000, and every 25th value was used for standard error calculation). Focusing first on similarities, under all three formulations, whether p is release-specific or a function of year, the covariates with the largest t -statistics are the site indicator for release from Sacramento, (log) flow, salinity, release temperature, exports and turbidity. The effect on survival of releasing at Sacramento, compared to further downstream at Courtland or Ryde, is a lowering of survival, as would be expected. Increases in flow, and salinity, are associated with increases in survival, while increases in release temperature have the opposite association. The adverse effect of water temperature increases on the survival of outmigrating juvenile chinook salmon in the Sacramento River was also reported by Baker *et al.* (1995). Of special interest to managers, the exports effect is negative under all three formulations; the effect appears statistically significant (using t -statistics) under the TBP and

Table 3 Estimated coefficients, and standard errors as subscripts, for models of S and p under the TBP, PL and hierarchical formulations

	Release-specific p			$p=f(\text{year})$		
	TBP	PL	Hier.	TBP	PL	Hier.
Intercept	1.31 _{0.06}	1.66 _{0.37}	0.59 _{0.10}	1.38 _{0.05}	1.75 _{0.32}	0.65 _{0.10}
Sacramento	-0.68 _{0.07}	-0.79 _{0.65}	-0.56 _{0.16}	-0.82 _{0.07}	-1.54 _{0.42}	-0.62 _{0.17}
Courtland	0.01 _{0.07}	0.31 _{0.57}	-0.02 _{0.17}	-0.17 _{0.06}	-1.89 _{0.41}	-0.09 _{0.17}
Size	-0.05 _{0.03}	-0.16 _{0.18}	0.23 _{0.06}	-0.10 _{0.02}	-0.32 _{0.13}	0.24 _{0.06}
Log flow	1.40 _{0.05}	1.63 _{0.38}	0.86 _{0.12}	1.38 _{0.05}	1.28 _{0.33}	0.86 _{0.12}
Salinity	0.53 _{0.03}	0.54 _{0.21}	0.30 _{0.09}	0.60 _{0.03}	0.66 _{0.15}	0.35 _{0.09}
Release temp.	-0.58 _{0.03}	-0.71 _{0.20}	-0.80 _{0.09}	-0.63 _{0.03}	-0.80 _{0.17}	-0.81 _{0.09}
Hatchery temp.	-0.31 _{0.03}	-0.37 _{0.20}	0.00 _{0.09}	-0.33 _{0.03}	-0.27 _{0.16}	-0.01 _{0.09}
Tide	0.09 _{0.02}	0.16 _{0.20}	-0.04 _{0.06}	-0.05 _{0.02}	-0.11 _{0.13}	-0.03 _{0.06}
Exports	-0.44 _{0.03}	-0.38 _{0.25}	-0.31 _{0.10}	-0.44 _{0.03}	-0.26 _{0.19}	-0.32 _{0.09}
Gate	-0.77 _{0.08}	-1.19 _{0.65}	-0.78 _{0.15}	-0.61 _{0.07}	-0.23 _{0.39}	-0.75 _{0.15}
Turbidity	1.33 _{0.05}	1.62 _{0.32}	0.38 _{0.13}	1.44 _{0.05}	1.64 _{0.26}	0.37 _{0.13}
Intercept p	—	—	—	-6.75 _{0.02}	-6.74 _{0.08}	-6.64 _{0.05}
1988 $_p$	—	—	—	0.70 _{0.04}	0.61 _{0.14}	0.60 _{0.09}
σ_S^2	—	—	0.14 _{0.01}	—	—	0.14 _{0.01}
σ_p^2	—	—	—	—	—	0.05 _{0.01}

Standard errors for the hierarchical model are calculated from thinned MCMC output. The covariates labelled Exports and Gate are interactions with indicators for release from Sacramento or Courtland.

hierarchical formulation, but not for the PL case. Similarly, the cross-channel gate being open has a negative effect and is statistically significant for the TBP and hierarchical cases but not the PL case.

The standard errors under the PL formulation are considerably larger than the TBP model, roughly five- to eight-fold increases, while the standard errors for the hierarchical model fell between the other two. This is as would be expected assuming overdispersion. Under the TBP model and using a t -test, all the covariates would be likely to be considered statistically significant, with the exception of the Courtland indicator.

The effect of modelling p on estimates of coefficients and standard errors for S was negligible for the TBP and hierarchical formulations. For the PL model, coefficients did change considerably for some covariates (e.g., Courtland indicator, size, flow) and the standard errors were consistently smaller for the case of p being a function of year. The estimated overdispersion parameters changed considerably; with release-specific p , $\hat{\phi}_{uo} = 84$ ($\hat{\phi}_{ut}$ was bounded below by 1.0), and with covariate based p , $\hat{\phi}_{uo} = 104$, with $\hat{\phi}_{ut} = 9$. The increased overall precision under the covariate-based p model, despite the increase in dispersion parameter values, is partially a reflection of the decrease in number of parameters to estimate (61 parameters for release-specific p vs. 2 for the covariate-based p).

The posterior distributions of the coefficients from the hierarchical model provide additional information about the relationship between estimated survival and the covariates. Figure 1 contains histograms of the posterior distributions for the survival coefficients. The percentage of values less than zero is shown above the histograms. Coefficients that are consistently above or below zero suggest a significant covariate effect on survival. With the exception of the Courtland indicator,

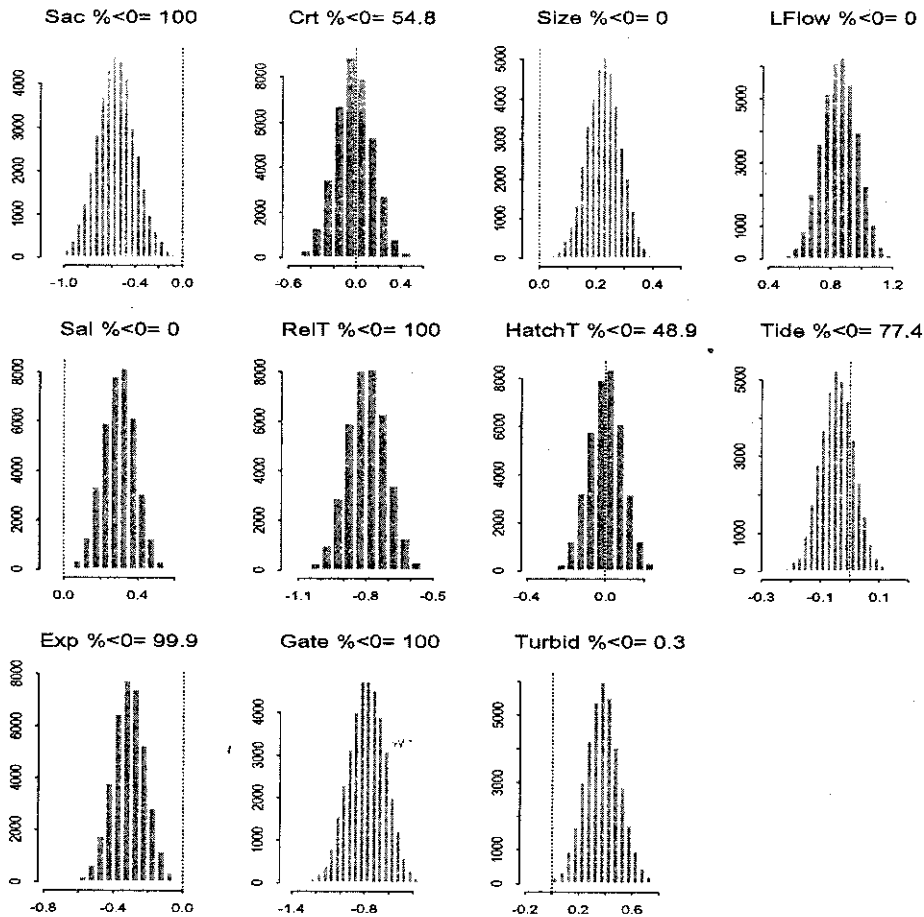


Figure 1 Posterior distributions of the coefficients for survival from the hierarchical model with release-specific p . Vertical lines are drawn at zero and the percentage of values less than zero are shown above the histograms

hatchery temperature and the tide variable, all the coefficients are strongly positive or negative.

The posterior means for σ_S^2 and σ_p^2 provide a measure of the magnitude of the release-specific random effects on S and p , respectively. Not surprisingly, given the difference in magnitude of S and p , σ_S^2 is considerably larger than σ_p^2 . Because the random effects enter into the calculation of S and p in a nonlinear way, it is simpler to look at the differences in what S would be if there was no random error, ϵ_S , in the logit link. An indirect measure is to compare $\exp(x'\beta)/(1 + \exp(x'\beta))$ to the posterior sample values for S , and similarly for p . This was done on a per simulated set of values, thus controlling somewhat for the uncertainty in the parameter estimates. The standard deviation of the difference between the estimated expected S , without random effects, and a random S was on average 3%, while for p the standard deviation was 0.008%.

4.2 Cross-validation

To compare the predictive ability and stability of the three procedures, a three-fold cross-validation was carried out. The cross-validation was restricted in that the training data subsets were chosen from the upstream releases alone and stratified samples were drawn with year of release being strata. The number of upstream releases in the three training subsets were 43, 40 and 40, while all 19 downstream releases were used. The restricted sampling was done to ensure that the complete set of ocean recovery probabilities could be estimated in all cases. Relatedly, only the covariate-based p model was fit because of the difficulty of doing predictions when the training set ps would necessarily differ from those for the test sets. Predictions were only made for trawl recoveries.

The coefficients for each method and in each training set are given in Table 4. For covariates which in the complete data set had significant coefficients under all three models, namely flow, salinity, release temperature and turbidity, the coefficients remained relatively large in absolute value and with the same signs. The Sacramento indicator was one exception for the TBP model with second training set. The hierarchical model was far more stable in the estimates, as the standard deviations of estimated coefficients were considerably smaller in most cases.

The predictive ability did not differ very much, however, between the three methods. Each method had a lower average absolute prediction error for exactly one of the three test sets.

5 Discussion

5.1 Sensitivity analysis

The assumption that the ocean recovery rate, π , is the same for a particular upstream and downstream release set is critical to the validity of the analysis, especially with regard to being able to separate survival and trawl capture probabilities. There are several reasons for this assumption to be wrong.

One reason is that downstream releases may experience some near immediate mortality after being transferred from the transporting truck and entering the river. This could be due to fatal temperature differentials, or disorientation, that make the fish more vulnerable to predators. Thus π would be lower for the downstream releases. Newman and Rice (1998) found some evidence of a shock, temperature differential effect on the upstream releases, in particular that there was a threshold level beyond which recovery rates for upstream releases worsened. Comparison of the release temperatures for the upstream and downstream members of a paired release, however, revealed that the water temperatures for the downstream releases were usually lower than those experienced by the upstream releases when they entered the river. This suggests that the shock effect, if present, might not have been as severe for downstream releases.

A related reason is that a culling of the weaker fish took place amongst the upstream releases and that those surviving the downstream migration and not being caught were

Table 4 Cross-validation results

	TBP			PL			Hierarchical			Coefficients' standard deviation ^a		
	Train 1	Train 2	Train 3	Train 1	Train 2	Train 3	Train 1	Train 2	Train 3	TBP	PL	Hier.
	Intercept	1.91	1.38	1.33	2.78	1.47	1.73	0.74	0.75	0.71	0.32	0.70
Sacramento	-1.07	0.15	-0.95	-2.14	-0.79	-1.38	-0.51	-1.03	-0.53	0.67	0.67	0.29
Courtland	-0.25	0.35	-0.45	-1.59	-0.93	-0.94	0.10	-0.51	-0.21	0.42	0.38	0.31
Size	-0.25	-1.27	-0.06	-0.55	-0.17	-0.31	0.15	0.00	0.42	0.10	0.19	0.22
Log flow	1.96	1.04	1.40	1.10	0.59	1.56	1.01	0.86	1.00	0.46	0.49	0.08
Salinity	0.78	0.36	0.51	1.02	0.23	0.69	0.32	0.27	0.49	0.22	0.40	0.12
Release temp.	-0.72	-0.65	-0.66	-1.49	-0.48	-0.89	-0.65	-0.83	-0.94	0.04	0.51	0.15
Hatchery temp.	-0.21	-0.37	-0.20	0.44	-0.38	-0.17	0.00	-0.15	-0.06	0.09	0.42	0.08
Tide	-0.06	0.21	-0.02	-0.19	-0.02	-0.03	0.07	-0.20	-0.09	0.14	0.10	0.14
Exports	-0.77	0.34	-0.58	-0.97	0.52	-0.50	-0.46	-0.03	-0.49	0.53	0.76	0.26
Gate	-0.46	-2.05	-0.43	0.30	-1.54	-0.26	-0.75	-0.51	-0.74	0.92	0.94	0.14
Turbidity	1.66	1.78	1.04	0.26	1.98	1.19	0.47	0.86	0.10	0.40	0.70	0.38
Intercept _p	-6.80	-6.73	-6.70	-6.88	-6.66	-6.69	-6.69	-6.60	-6.62	0.05	0.12	0.05
1988 _p	0.65	0.54	0.76	0.68	0.50	0.62	0.50	0.60	0.68	0.11	0.09	0.09
PE	20.23	20.38	16.26	21.56	22.37	13.23	21.74	18.82	19.26			

Estimated coefficients, for covariate-based *p* model, for the three methods based on three training data subsets of the upstream releases. Values reported for the Hierarchical model are the means of the posterior distributions. |PE| is the average absolute prediction error for the corresponding test sets. ^aBetween-training subset coefficients' standard deviations.

a group of relatively strong fish compared to the more heterogenous downstream release group.

Conversely, upstream releases could experience a delayed mortality due to some factor encountered between the point of release and the in-river trawl. In other words they survived to the in-river trawl and eluded capture but something upstream of the trawl has fatally harmed some of them. This would cause π to be lower for the upstream survivors than for the downstream releases, but seems less tenable than the above situations.

A fourth reason for differing π s is differences in the ocean distribution and migration pattern between upstream and downstream releases. Because of the spatial and temporal irregularity of the ocean fishery, this could lead to different harvest rates on the two groups.

Lastly, genetic and rearing environmental differences could exist between upstream and downstream releases. The upstream releases always came from a single hatchery, Feather River Hatchery, but the downstream releases came from one of two hatcheries, Feather River Hatchery or Coleman National Fish Hatchery.

Assuming that downstream releases experience mortality higher than for upstream releases reaching the downstream location, and that this is expressed as a fixed multiple of the mortality rates of surviving upstream releases, the TBP model can be extended as follows:

$$\hat{y}_{do}|\pi, \psi \sim \text{Binomial}(R_d, \psi\pi) \quad (5.1)$$

where $0 \leq \psi \leq 1$. Extensions for the PL and hierarchical models are similar. Given multiple release sets of one or more upstream releases paired with single downstream releases, the 'shock' effect, ψ , is estimable assuming that S can be modelled as a function of covariates. The TBP, PL and hierarchical models were re-fit with the shock parameter ψ . The prior for ψ for the hierarchical model was uniform on $(0,1)$.

The results for the models including ψ are summarized in Table 5. The estimates of the shock parameter, ψ , differed considerably between the model formulations, with the hierarchical model indicating a much stronger effect than the TBP and PL models. The magnitude of the variance of the logistic normal error term for survival, σ_S^2 , decreased nearly 30%.

With the addition of the ψ parameter, the estimated values for S went down while estimates of p and π increased. For example, the average π for the TBP model increased 48% for the release specific p case. Note, in particular, the difference in the intercepts in Tables 3 and 5. Such changes are not unexpected given that \hat{y}_{do}/R_d essentially estimates π in the without shock case and the product $\psi\pi$ in the shock case. If shock is present, that is, $\psi < 1$, and assuming the quality of the fit to \hat{y}_{do} will not worsen by adding another parameter, π must increase. And if π increases, the products Sp and $S(1-p)$ will likely decrease to maintain the same quality of fit to y_{ut} and \hat{y}_{uo} . Relatedly, presence of a shock effect in downstream releases suggests the presence of a shock effect in upstream releases, too, but the effect is absorbed by S . The use of estimates of S based on experimental releases as proxies for the survival of naturally outmigrating juveniles could be overly pessimistic given the latter remain in the same water.

Table 5 Sensitivity analysis

	Release-specific p^a			$p = f(\text{year})^a$		
	TBP	PL	Hier.	TBP	PL	Hier.
Intercept	0.27 _{0.07}	0.47 _{0.34}	-1.02 _{0.08}	0.75 _{0.07}	1.16 _{0.37}	-1.02 _{0.10}
Sacramento	-0.72 _{0.05}	-0.96 _{0.33}	-0.25 _{0.13}	-0.85 _{0.05}	-1.38 _{0.34}	-0.24 _{0.13}
Courtland	-0.15 _{0.04}	-0.09 _{0.30}	0.11 _{0.14}	-0.26 _{0.04}	-0.74 _{0.33}	0.04 _{0.13}
Size	-0.02 _{0.02}	-0.04 _{0.11}	0.11 _{0.05}	-0.06 _{0.02}	-0.22 _{0.12}	0.10 _{0.05}
Log flow	1.15 _{0.04}	1.25 _{0.26}	0.57 _{0.09}	1.22 _{0.04}	1.17 _{0.26}	0.56 _{0.09}
Salinity	0.46 _{0.02}	0.44 _{0.14}	0.21 _{0.07}	0.56 _{0.03}	0.59 _{0.13}	0.23 _{0.07}
Release temp.	-0.44 _{0.02}	-0.47 _{0.14}	-0.55 _{0.07}	-0.54 _{0.02}	-0.66 _{0.15}	-0.56 _{0.07}
Hatchery temp.	-0.31 _{0.02}	-0.33 _{0.14}	-0.12 _{0.07}	-0.32 _{0.02}	-0.28 _{0.13}	-0.10 _{0.07}
Tide	0.03 _{0.02}	0.07 _{0.11}	-0.02 _{0.05}	0.03 _{0.02}	-0.08 _{0.11}	0.00 _{0.05}
Exports	-0.39 _{0.03}	-0.35 _{0.16}	-0.20 _{0.07}	-0.42 _{0.03}	-0.27 _{0.16}	-0.21 _{0.07}
Gate	-0.40 _{0.04}	-0.47 _{0.31}	-0.63 _{0.12}	-0.38 _{0.04}	-0.19 _{0.30}	-0.60 _{0.13}
Turbidity	1.08 _{0.04}	1.16 _{0.23}	-0.21 _{0.12}	1.26 _{0.04}	1.43 _{0.23}	0.04 _{0.10}
Shock	0.64 _{0.02}	0.65 _{0.07}	0.40 _{0.01}	0.76 _{0.02}	0.80 _{0.09}	0.41 _{0.01}
Intercept p	—	—	—	-6.61 _{0.03}	-6.62 _{0.11}	-5.93 _{0.06}
1988 p	—	—	—	0.80 _{0.05}	0.65 _{0.15}	0.68 _{0.09}
σ_S^2	—	—	0.10 _{0.01}	—	—	0.10 _{0.01}
σ_p	—	—	—	—	—	0.06 _{0.01}

Estimated coefficients, and standard errors as subscripts, for models of S (and covariate based p) under the TBP, PL, and hierarchical formulations when temperature shock (Shock) is allowed for downstream releases. The covariates labelled Exports and Gate are interactions with indicators for release from Sacramento or Courtland. ^aIncluding shock effect for downstream releases.

General conclusions regarding which covariates appear to be most influential on S and the nature of their relationship with S do not change, however, from the models that excluded shock. Excluding the intercepts, the correlations between the t -statistics for the β s exceed 0.97. Based on the hierarchical model with covariate-based p , for example, the coefficients with posterior distributions largely completely positive or negative were the Sacramento indicator (97% negative), size (99% positive), flow and salinity (both 100% positive), release and hatchery temperatures (100 and 94% negative), exports and the cross-channel gate indicator (both 100% negative). The inclusion of hatchery temperature in this set is the sole contrast with the case without ψ .

5.2 More complex hierarchical models

The hierarchical formulation used is arguably an improvement over the pseudo-likelihood approach in that the latter can be viewed as an approximation to a hierarchical model for the parameter combinations Sp and $S(1-p)\pi$ of the TBP. The hierarchical model separated the random effects of S and p in particular, which seems more realistic than modelling the parameter combinations.

One assumption of the hierarchical formulation, however, is that for a given release group there is just one set of S , p and π values. Thus the variation in (S, p, π) combinations is just between releases. With releases numbering 50 000 or so fish, there is undoubtedly variation in (S, p, π) within releases, too. It can take one to two hours to release the fish into the river and variation in the subsequent downstream path

Henry Stewart Conference Studies

Forthcoming Statistics Events

October 2003

***Current Thinking on The Use of Bayesian Analysis in the
Pharmaceutical Industry – European Event***

Tuesday, 14th October 2003,
London, UK

***Current Thinking on The Use of Bayesian Analysis in the
Pharmaceutical Industry – U.S. Event***

Wednesday, 22nd October 2003,
Washington DC, USA

Statistics of Optimal Dosing

Wednesday, 28th October 2003,
Washington DC, USA

November 2003

***Solving the Regulatory and Statistical Issues Associated with the Use of
Control Groups***

Wednesday, 22nd October 2003, Washington DC, USA

Full details of all the speakers and topics to be covered can be found on the Henry Stewart website www.henrystewart.com

For European events please contact Dr Carlos Horkan, Tel: +44 207 404 3040; Fax: +44 207 404 2081; Email: CarlosH@henrystewart.co.uk

For U.S. events please contact Vivien Maier, Tel: 1-800 229 2416; Fax: 206 244 2681 or by Email: vivienmaier@vmainc.com

Full details can be found on the Henry Stewart website www.henrystewart.com or alternatively by contacting Dr Carlos Horkan, Tel: +44 207 404 3040; Fax: +44 207 404 2081; Email: CarlosH@henrystewart.co.uk