November 2015

Estimation Problems in Complex Field Studies with Deep Interactions: Time-to-Event and Local Regression Models for Environmental Effects on Vital Rates

Krzysztof M. Sakrejda
University of Massachusetts - Amherst

Estimation Problems in Complex Field Studies with Deep Interactions:
Time-to-Event and Local Regression Models for
Environmental Effects on Vital Rates

A Dissertation Presented
by
Krzysztof M. Sakrejda

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

September 2015

Organismic and Evolutionary Biology Program

Estimation Problems in Complex Field Studies with Deep Interactions:
Time-to-Event and Local Regression Models for
Environmental Effects on Vital Rates

A Dissertation Presented
by
Krzysztof M. Sakrejda

Approved as to style and content by:

_____

Benjamin H. Letcher, Chair

_____

Michael L. Lavine, Member

_____

Adam H. Porter, Member

_____

Ethan D. Clotfelter, Member

_____

Elizabeth R. Dumont, Director
Organismic and Evolutionary Biology

To Dr. Janina Morgiel and Dr. Iwona Sakrejda,
who motivated me to start down this path.

To Dr. Alexandria Brown, who motivated me to finish it.

To Rowan and Stella, the children who learned that
"pretty please with a finished dissertation on top"
was the best way to get ice cream.

Abstract

Estimation Problems in Complex Field Studies
with Deep Interactions:
Time-to-Event and Local Regression Models for
Environmental Effects on Vital Rates

September 2015

Krzysztof M. Sakrejda, B.A. Vassar College

M.S. University of Oregon

Ph.D. University of Massachusetts, Amherst

Directed By: Dr. Benjamin H. Letcher

Field studies that measure vital rates in context over extended time periods are a cornerstone of our understanding of population processes. These studies inform us about the relationship between biological process and environmental noise in an irreplaceable way. These data sets bring "big data" and "big model" challenges, which limit the application of standard software (e.g., **BUGS**). The environmental sensitivity of vital rates is also expected to exhibit interactions and non-linearity, which typically result in difficult model selection questions in large data sets. Finally, long-term ecological data sets often contain complex temporal structure. In commonly applied discrete-time models complex temporal structure forces the analyst to make many competing decisions about discretizing time, which add complexity to any analysis and confuse the interpretation of parameters. I tackle these problems in three ways: 1) I apply and improve current state of the art statistical algorithms and software (Hamiltonian Monte Carlo with the NUTS sampler implemented in **Stan**) to implement the estimation of vital rates from mark-recapture data in a large Atlantic salmon data set; 2) I extend continuous-time time-to-event models with continuous estimated rate functions to mark-recapture and emigration data; and 3) I apply simple local regression to the estimation of vital rates to turn the model-selection problem resulting from interactions and non-linearity into an estimation problem, which is solved by the estimation algorithm. For each model I present methods for visualizing the results in biologically meaningful terms and use these for model criticism.

Contents

List of tables

LIST OF FIGURES

x

CHAPTER 1

VITAL RATES AS AN ESTIMATION PROBLEM

Constructing relationships between vital rates and environmental covariates is an important
current topic due to increasing demand for predictions for a variety of ecological systems.
While complex specialized models have been created for some systems, their development
requires advanced statistical knowledge. I see a recurring set of requirements that have stim-
ulated the development of complex models: 1) flexibility in relating the data to a process
model and often also to an observation model; 2) flexibility in the timing of observations;
3) flexibility in discovering and estimating non-linear responses of process to covariates and
combinations of covariates; 4) flexibility in discovering and estimating seasonal variation in
the response of process to covariates; 5) an ability to explore the quality of the estimates;
and 6) a method for evaluating data requirements in the context of the desired quality
of estimates. For some systems these requirements do force ecologists to rely on special-
ized models. For an important subset of models I suggest that a general set of methods
can be constructed from available, if advanced, statistical tools that addresses all of these
requirements.

My interest in this problem arises from the analysis of a long-term mark-recapture study
of fish in a small stream system that resulted in 21,386 captures of tagged Atlantic salmon
over 47 recapture occasions. The Atlantic salmon data set spans more than a decade and
data collection continues for other species at the site. While previous analyses have used
this data set to estimate how survival of Atlantic salmon is affected by seasonality and age
structure (B. H. Letcher, Dubreuil, *et al.*, 2004; B. H. Letcher & Horton, 2008; B. H. Letcher,
Gries, *et al.*, 2002), I focus primarily on the response to environmental covariates.

The salmon study displays variation in timing of sampling common to complex field
studies. Sampling occasions for this study are scheduled around early April, early June,
mid September, and early December based on salmon biology. Carrying out a sampling
occasion requires two weeks of intensive in-stream work by a team of people that can be
impossible to complete during storms and potentially for weeks after extensive rainfall. For
this reason and due to other logistical constraints, samples during the decade of the study
are spread over much of the year.

Field conditions from the salmon study also challenge our ideas of environmental ef-

1

fects. Environmental tolerances of Atlantic salmon are well defined in the laboratory, and field measurements of environmental effects seem almost superfluous. However, according to laboratory measurements, stream temperatures at the site rarely enter the zone that prevents feeding in Atlantic salmon. During the decade of the study stream temperatures never entered the zone that would cause mortality. Salmon are clearly capable of coping physiologically with the range of conditions available at the West brook.

A straightforward hypothesis would be that any environmental effects on survival should be either weak or non-existent. If the animal can survive these conditions confined to a tank, it can surely survive just as well in its natural habitat. An alternative hypothesis would be that in the field environmental effects act through mediating mechanisms. Low water conditions might deplete habitat, increase antagonistic interactions, and make the remaining stream habitat more accessible to predators. Such indirect effects could clearly affect survival even though stream flows would never enter a range that posed a physical or physiological challenge to salmon. Though indirect effects have been demonstrated (Breau et al., 2007; Cunjak et al., 1998), little is known about the range of conditions that cause them to become relevant for survival.

Adding another layer of complexity to the question of environmental effects, experience at the study as well as the physiological and behavioral literature tell us that salmon react to the strong seasonality of the study system. They seek season-specific habitats, season-specific food intake rates, and are more or less territorial according to the time of the year. I expect that environmental responses might also vary seasonally. This issue is critical because it cuts roughly by four the sample size for estimating environmental response surfaces, leaving about 11 sampling occasions per season.

The tools for addressing these issues are, broadly speaking, available in the statistical literature. The challenge in applying them to ecological systems is that naive applications do not meet the requirements outlined in the opening paragraph above, and the tools are generally not appropriate for complex field data. Making them applicable requires non-trivial modifications.

This dissertation presents a series of methods for discovering relationships between vital rates and environmental covariates using field data. I discuss the likelihood calculations required to fit each model along with the details required to extend calculations to related models. I present each method with a simple illustrative example that includes a statistical interpretation but limited biological interpretation. Throughout the dissertation I return to

2

the Atlantic salmon system as both a statistical example and in order to answer biological questions about the sensitivity of this species to key environmental covariates.

First I treat Atlantic salmon survival in detail, considering a set of models based on the Cormack-Jolly-Seber mark-recapture model. Rather than replacing this basic model I modify it to incorporate some ideas from time-to-event modeling in order to simplify the calculations for this large data set. I use polynomial regression models with interactions to produce response surfaces describing the seasonal sensitivity of survival to stream flow and water temperature. While these surfaces are plausible they are difficult to evaluate at given points in covariate space for their fidelity to the data.

Second, I illustrate how time-to-event models can be applied to univariate and multivariate event data collected under common observation schemes with and without uncertainty. Time-to-event data includes data on emigration and survival that is critical for measuring vital rates, as well as other studies where the time until a discrete outcome is measured. The time-to-event literature is particularly rich and spans multiple fields. Rather than attempt a comprehensive treatment I focus on inhomogeneous Poisson process (IHPP) models with both continuous and piecewise constant intensity functions. These models have a clear mathematical foundation and are appropriate for many problems with time-varying event rates driven by outside forces. I apply these model to a laboratory example as well as our Atlantic salmon data set.

Third, I introduce local regression as a method parameterizing the environmental surfaces in a more flexible way. Local regression makes it easier to designate which parts of the estimated surface are reliable by relating both data and parameters to specific parts of the covariate space. The method is based on seasonal multi-dimensional splines linked to arbitrary likelihood functions using the moderately user-friendly and very extensible Stan modeling language. Estimation is handled through a variety of algorithms included in Stan, and the trade-off between data requirements and flexibility of the associated relationships is tuned through the definition of knot points. I include methods for visualizing the resulting relationships and for evaluating uncertainty of estimates, as well as implications for data requirements. The goal is to provide a model that could be used throughout a field study to adapt study design and maximize efficiency of data collection.

1.1 Prediction context

Climate change is expected to drive changes in species distributions in the future (Peterson *et al.*, 2002) as it has in the past (Chen *et al.*, 2011; D'Andrea *et al.*, 2011). These spatial species shifts will impact the availability of species for economic exploitation (A. H. Hines *et al.*, 2010), as well as the effectiveness of conservation strategies particularly for species on high-elevation islands of cool habitat (Dullinger *et al.*, 2012), freshwater fish (Lyons *et al.*, 2010; Xenopoulos *et al.*, 2005), and ice-dependent circumpolar species (Lynch *et al.*, 2012). Our ability to respond to species declines, expansions, and invasions is dependent on our ability to make accurate predictions about potential future shifts in species distributions and the processes that will bring them about. An interest in prediction is not new to ecology but implementing practices to achieve the promised benefits (J. S. Clark *et al.*, 2001) has proven difficult. Our contribution focuses on estimation in the prediction context, but I first expand on that context here.

Ecology possesses a number of frameworks that have been applied to prediction problems. Matrix projection models (Caswell, 2001) have a rich theory and are rooted in the field of population biology. These models were historically restricted to discrete classification of population sub-groups but in recent years they have been extended to Integral Projection Models (IPMs) in order to address continuous population parameters (Ellner & Rees, 2006) while keeping much of the useful theoretical machinery. In the IPM framework the focus is on vital rates—reproductive rates, death rates, and emigration/immigration as well as movement rates. These models are expressive and attractive exactly because they directly model the parameters (vital rates) that capture population-level dynamics and can be used to predict future population trajectories.

In contrast with the population-level approach of IPM's, Individual Based Models (IBM's) typically focus on parameterizing individual biology based on trait values and environmental responses. Starting from a given population, growth and reproductive histories of individuals are simulated forward. Finally, to characterize the possible future trajectories of the population, such individual histories are reduced to summary statistics such as population size and growth rate. Due to the focus on individual level parameters, IBM's can benefit from the detailed literature on physiological and behavioral responses of animals.

Tests of predictions based on IPM's and IBM's are important to characterizing the state of the art with regards to prediction in ecology. However, such tests are also rare, likely

because the models are data-intensive and data collection is expensive and time consuming relative to other methods. Crone *et al.* (2013) show that the state of the art is weak and provide some insight into the mechanisms. They constructed stage-specific matrix models for 82 plant populations and 20 species using population-specific methods and used them for deterministic as well as stochastic predictions of population sizes. These matrix models were constructed from relatively high quality data—primarily stage-specific observations of individual fates. While the matrix models did effectively represent the data (good in-sample predictions), they performed very poorly in a five-year out-of-sample prediction exercise.

The conclusions of Crone *et al.* (2013) included two key points about ecological prediction: 1) for poorly studied systems we can not expect predictive models to perform accurately; and 2) even for well-studied systems, understanding responses to environmental covariates is critical for accurate predictions. We acknowledge that many important systems are not well studied, and that the development of techniques for prediction in these systems is a critical area of research(Contamin & Ellison, 2009; Hastings & Wysham, 2010). Our focus in the present manuscript is on species that have attracted sufficient resources to be well-studied and yet remain poorly characterized in the context of quantitative ecological predictions. Such species include many species of economic importance(Brodeur & Wilson, 1996; Krkosek & Drake, 2014), scientific model systems that could be used to understand prediction problems in more generality, as well as other species that have attracted attention for a variety of reasons including conservation status and cultural importance.

Constructing broadly applicable estimates of environmental effects is not a purely statistical problem as it combines statistics along with the *desired interpretations of effects relevant for prediction*. Estimates embed assumptions and when those assumptions are not met in prediction we should expect poor results. For example, estimating individual-level environmental effects under laboratory conditions may not be relevant for predicting the fraction of surviving individuals in the field under a given set of conditions. The general topic is broad and not united in the literature. This dissertation contributes by focusing on estimation of environmental responses in the context of deep interactions from complex field studies.

Although environmental effects are clearly important in some form, there are many ways environmental effects could be included in predictive models. The paucity of out-of-sample model testing and validation exercises in the IPM/IBM context means we are left with no guidelines on which environmental effects to pursue, how to include them, and what level

of detail is relevant. Conveniently, this question is addressed extensively in the literature on species distribution models—a context where data collection, estimation, and prediction are all simpler exercises.

1.2 Species distribution models

Species distribution models have been broadly applied to answer questions about the potential impact of climate change. Species distribution models (SDMs) are sometimes called environmental niche models, bio-climatic envelope models, or environmental envelope models depending on the focus of the analyst. The typical application of an SDM is to describe the realized niche of an organism and its relationship to biotic and climatic predictors (Guisan & Thuiller, 2005). The primary goal of model building is to describe correlations that are stable enough to provide ecological insight or generate accurate spatial and temporal predictions (Guisan & Zimmermann, 2000; Wenger & Olden, 2012). Environmental niche models are fitted to data using a variety of techniques from statistics and machine learning including generalized linear mixed models, artificial neural networks, generalized additive models, and maximum entropy (Downie *et al.*, 2013; Guisan & Zimmermann, 2000; Sundblad *et al.*, 2009; Wenger & Olden, 2012). Different computational choices can affect inference but they are typically less important than overall questions about the meaning of the inferred relationships.

The early interest in SDM's (Guisan & Zimmermann, 2000) and more recent expansion in their use has been driven in part by their low data requirements and access to cheap GIS- or satellite-derived predictors. The cost advantage has allowed researchers to apply SDM's on a continental scale (Logez *et al.*, 2012), which would be very difficult both in terms of data requirements and computationally for a more dynamic model.

Due to the confusing range of predictors available in the SDM context, some effort has been put towards classifying predictors. Some predictors, particularly abiotic predictors, can be accurately measured remotely making large-scale analysis less costly (Snickars *et al.*, 2014). Other predictors, particularly biotic predictors, are more costly to measure and if used they make models more difficult to apply for spatial prediction. These practical aspects of predictor choice must be weighed against theoretical concerns. Specifically, Guisan & Zimmermann (2000) advocates the classification of predictors (and their gradients) as *resource*, *direct*, and *indirect*. Resources are predictors that, when used by an organism, are

not available to others, such as territories, food, or insolation. Direct predictors affect the organism directly but are not consumed in the process such as water temperature or water velocity for a stream fish. Indirect predictors have no direct effect on the organism but are correlated enough with direct predictors to be useful. The correlations between indirect predictors and more direct predictors are not necessarily stable in space or time but their low cost and wide availability have made them a staple of species distribution models.

Species distribution models have widely acknowledged caveats in their interpretation due to their reliance on correlations (Guisan & Zimmermann, 2000). The main caveat is that they measure the realized niche whereas their application to prediction would be better served by describing the fundamental habitat niche, which may be larger or smaller. Common sources of mismatch are dispersal, interspecific competition or mutualisms, and non-equilibrium populations either used in model fitting or in model validation. Due to these caveats some recent modeling efforts have focused on the spatial and temporal predictive ability of species distribution models. We discuss both spatial and temporal predictive ability as "transferability" in the sense that the goal is to train a model on one data set and transfer the inferred responses to environmental variables to another setting separated in space and/or time.

Focusing on the ability to predict species presence or absence, Randin *et al.* (2006) assessed bi-directional transferability between a study area in Switzerland and one in Austria. They assessed transferability for SDM's describing 54 plant species based on topography and climate variables. They found that transferability varied widely by species and that the models were not transferable for approximately 50-70% of species. The index used to assess transferability is discussed in detail by Randin *et al.* (2006). Sundblad *et al.* (2009) found better overall results for two fish (northern pike and roach) but also showed that transferability could be directional, which highlights the need to understand the relationship between sites used to estimate a model and its the sites it will be applied to (Barbet-Massin *et al.*, 2010).

Ultimately it is not environmental correlations but vital rates that control species range distributions (Guisan & Zimmermann, 2000). In line with the focus on vital rates, Austin (2002) pointed out that predictors can be thought of as variables chained together by correlations with a distal end and a proximal end relative to vital rates. The very distal end moves freely and it encompasses variables, such as longitude, that are only loosely related to conditions experienced by the organism. The proximal end might include ground-level

insolation for a forb or the body temperature of an ectotherm. The proximal end is directly tied to the biological process affecting the organism. While distal variables are easier to measure we rely on ecological theory and data to establish how to best relate them to proximal predictors and ultimately vital rates. In toto, this body of work suggests that assessing direct variables and their relationship to vital rates is an appropriate area to explore in order to advance our ability to make generalizable predictions about climate-driven changes in species distributions.

From a theoretical point of view it is expected that predictors relevant to vital rates would be best for prediction. However, this expectation has not been thoroughly tested in the matrix model or integral projection model literature with the exception of (Crone *et al.*, 2013). It is the species distribution modeling literature, with its cheaper estimation, prediction, and validation framework that convincingly demonstrates the importance of vital rates.

Frequently the relationship between proximal variables and biological responses is assessed in the laboratory in studies of animal physiology and behavior. For Atlantic salmon, our focal organism, this literature yields a rich collection of detailed and generally consistent measurements of salmon response to temperature stresses and different flow regimes, as well as behavioral responses to starvation or predator exposure. From the perspective of individual based models it appears as though we have the required response estimates in hand for constructing a complete model. However, these responses are not necessarily directly applicable to prediction on larger temporal and spatial scales. First, lab studies and small scale field studies by definition do not measure endpoints relevant to long-term or large-scale prediction. For example a lab study might measure the response of 7-day survival to extreme temperatures whereas monthly or quarterly survival would be relevant. Second, a lab study by definition reduces variation by removing environmental influences that will operate in a field setting. Some of these might simply increase variation whereas others will bias the measured outcome (e.g., survival under temperature stress in the absence of predators will be higher than with predators present). I am not aware of a literature to support our assertions that laboratory studies are difficult to apply in the prediction context, but I do have estimates for a wide range of parameters available and I compare them to our environmental response surfaces for Atlantic salmon to qualitatively evaluate their utility.

1.3 Key challenges

The general issue of parameter estimation for prediction is very broad and impacts a variety of fields. Prediction exercises to compare competing approaches have sprung up in a number of fields from epidemiology to genetics. Rather than attempting a comprehensive treatment, I focus on a few important challenges.

In the first chapter I introduce our most complex field study and recurring example—an Atlantic salmon (*Salmo salar*) mark recapture data set collected between 1997 and 2008. I describe the entire set of field methods used to collect this data set as well as the relevant life history and physiological or behavioral literature on responses to key environmental parameters. This data set is used for three analyses: 1) an analysis of the time until emigration; 2) an analysis of survival based on a Cormack-Jolly-Seber model adapted to varying time intervals and multiple outcomes; and 3) a final-reanalysis of the survival data with the same model but using local regression to describe the relationship between survival and environmental covariates.

Long-term intensive field studies often include mark-recapture data, which is a cost-effective way of monitoring a population in detail, by marking a hopefully representative sub-population and monitoring its fate. We take one such data set and modify the Cormack-Jolly-Seber model to account for multiple known and unknown fates and varying time intervals between recaptures. This treatment is a half-step to a continuous-time model but it makes it possible to construct the environmental response model without also attempting a new model type.

Multi-year (or otherwise long-term) field studies are critical to evaluating population responses to climate change. They are a bridge between detailed laboratory studies and large scale species distribution models. Long-term field studies often result in complex temporal structure in the observations due to logistical constraints and technological changes. In contrast, most analytical methods for field data are designed around discrete time steps and complicate the analysis of field data. In the third chapter, I demonstrate how continuous-time methods for the analysis of discrete events can be applied to a variety of problems typically analyzed using discrete-time methods. I use two examples—the Atlantic salmon data set as well as a behavioral study with multi-variate outcomes.

In ecological systems we do not expect effects to act alone, but experiments that measure multiple interacting covariates are difficult to model. The quality of inference in such

complex models is difficult to evaluate. In the fourth chapter I re-use the survival model framework from the second chapter but modify it to use a flexible method for both quantifying complex interactions and evaluating how well the data support the estimated response surface. I illustrate this method with a small data set of stream temperature observations and the larger analysis of Atlantic salmon survival.

Natural history and field study

In the remainder of the dissertation I use a three data sets. I use measurements from one behavioral study on shiners and sunfish conducted by Fuller & Earley (2015). I avoid the discussion of the scientific motivation for this study as I only became involved in the somewhat complicated analysis after the experiment was carried out. I also use a small subset of a dataset on water temperatures from streams in New England (B. Letcher *et al.*, n.d.). The original motivation for this dissertation and the most complex examples I use are a subset of the mark-recapture data set collected by the Letcher lab since 1997 and currently extended by a larger group. In addition to the methods described in the following chapters, I also contribute an analysis of stream flow and temperature effects on Atlantic salmon. I describe the life history of Atlantic salmon and their environmental sensitivity here along with a compilation of the field study methods.

2.1 Salmon freshwater life history

Atlantic salmon are an iteroparous salmonid that begin their life in small freshwater streams. Generally they remain in their natal stream through the most of the juvenile portion of their life cycle before migrating out to the ocean to complete growth, mature, and return to the natal stream to spawn. The Atlantic salmon life cycle is a model system for life history variation and admits many departures from this general scheme. We present a detailed summary of Atlantic salmon life history here to provide a context for the interest in flow and temperature effects on survival. The life history of Atlantic salmon, and therefore salmons' interaction with the environment, is strongly structured by stage.

Atlantic salmon begin life as fertilized eggs, buried at a depth of 15-25 cm, in coarse gravel-to-cobble sized material. These sites, called redds, are found beneath fast ($0 - 50$ cm s^{-1}), shallow ($20 - 30$ cm) water (Bardonnet & Baglinière, 2000). Redds vary widely in their outward appearance and location but on a small scale they are unified by the eggs' need for oxygenated water (Rubin & Glimsater, 1996), as well as protection from predators and high-flow events. The development of the eggs is controlled by temperature but after a period of 4-5 months (Heggberget, 1988), the surviving eggs hatch. The resulting alevin

continue to develop and consume their yolk sac before they swim up through the gravel into their new habitat. The emerging fry are small ($\sim$ 25 mm) and with their yolk exhausted they turn to preying on benthic invertebrates (Petersson *et al.*, 1996) and searching for suitable habitat. The search for prey and shelter exposes young fry to strong inter- and intra-specific predation (Leaniz *et al.*, 2000) and makes them vulnerable to mortality from high flow events (Jensen & Johnsen, 1999). The limited endurance and swimming ability of small fry result in primarily downstream movement though dispersal distances can be significant (e.g.-$\sim$ 1 km, (Webb *et al.*, 2001)). Fry prefer riffles under 10 cm deep and slow flows ($20 - 50$ cm per second) for habitat through their first summer and transition to more open habitat in the first autumn following hatching (Rimmer *et al.*, 1984; Saunders & Gee, 1964).

In the autumn of their first year fry continue to grow, darken in colour, and develop vertical marks on their sides. When the marks are developed juvenile salmon are referred to as "parr". Atlantic salmon parr use deeper riffles ($25 - 50$ cm) and faster flowing water ($10 - 65+$) than most other freshwater stream fish as they are able to use their fins to hold position on the stream bottom without actively swimming (Arnold *et al.*, 1991). When parr settle they typically stay in their chosen home range (Saunders & Gee, 1964; Steingrímsson & Grant, 2003) over multiple seasons or years. If the salmon survive the winter, they end their young-of-year (YOY) stage and are referred to as over-yearling fish. This time point is important as it marks the beginning of the first high-growth spring period. Up through the young-of-year stage, the changing interaction between salmon and their environment is driven primarily by their changing size and ability to move through the environment.

For over-yearling fish, size is still important but the interaction with seasonality increases in importance. The age 1 winter following the YOY stage is the first season when growth is minimal due to a voluntary reduction of food intake (N. B. Metcalfe & Thorpe, 1992). Reduced food intake results in dropping body lipid reserves (Berg & Bremset, 1998) and it is mediated by a behavioral mechanism that triggers increased feeding when lipid reserves fall below a threshold (Finstad, Berg, *et al.*, 2010; Finstad, Ugedal, *et al.*, 2004). The winter season is typically a time of severely reduced activity (Cunjak, 1988; Gardiner & Geddes, 1980). Increased activity by fish with the smallest energy reserves appears to result in increased mortality, likely from increased exposure to predation (Berg & Bremset, 1998; Finstad, Ugedal, *et al.*, 2004). The pattern remains the same for older fish in the winter season (Finstad, Ugedal, *et al.*, 2004), except at least some older fish begin to prepare for

smolting and maintain growth through the winter (Sigourney, 2010).

Besides reduced activity and increased preference for shelter (Cunjak, 1988), over-wintering fish also face decreasing stream flows and habitat availability during the early winter period. During the mid-winter period anchor ice formation and melting affect habitat availability (Linnansaari *et al.*, 2009) whereas decreased water quality and habitat scouring become more important during the warming phase (Cunjak *et al.*, 1998). Following the winter, most growth for the year takes places in a short period during the spring season (Bacon *et al.*, 2005).

While spring snowmelt and the resulting flood conditions can rearrange the stream bed and present a challenge for many stream organisms, Atlantic salmon are uniquely adapted to maintain position without swimming (Arnold *et al.*, 1991) and appear to not be strongly affected (Enders *et al.*, 2005; Jensen & Johnsen, 1999). High flows during this period can have positive (Davidson *et al.*, 2010) or negative (Jensen & Johnsen, 1999) effects on growth and only weakly affect survival (Bacon *et al.*, 2005; Jensen & Johnsen, 1999) of over-yearling fish. Stream temperatures also change rapidly in spring and in laboratory conditions this often leads to increased growth (J. M. Elliott & Hurley, 1997). However, the reverse can be true in natural conditions (Davidson *et al.*, 2010) so the direction and magnitude of any temperature effects on growth are uncertain.

As stream flows decrease and temperatures climb in the summer months, heat stress becomes an important challenge for salmon (Breau *et al.*, 2007). The physiological limits for growth and survival in Atlantic salmon and related species are well characterized (J. M. Elliott & J. A. Elliott, 2010), but these physiological limits can not be directly compared to measured in-stream temperature due to a variety of behavioral mechanisms. Even in a stream where overall water temperature is above the physiological maximum, cold ground-water seeps or inputs from smaller, cooler streams create zones of acceptable environmental conditions termed "thermal refugia"(Torgersen *et al.*, 1999) and Atlantic salmon have been shown to exploit these (Breau *et al.*, 2007).

Early autumn marks a further decrease in growth from the summer (Finstad, Berg, *et al.*, 2010). Decreased growth is also matched by a shift towards slower water velocities that would have fewer feeding opportunities and better cover (Rimmer *et al.*, 1984). Stream-wide flows increases and temperatures drop making temperature and oxygen stress less relevant at this time of the year. This discussion of seasonality for over-yearling fish applies best to those that remain as parr, and some that will smolt late in the following year. Another group

of fish that will smolt early in the following spring continue to grow through the autumn and winter seasons, which suggests that they continue their regular foraging activities through the cool autumn and winter season.

2.2 Summary of known mechanisms affecting vital rates

Based on the Atlantic salmon life history, I divide effects on survival into four groups: 1) metabolic issues affecting time budget and predation; 2) shelter availability; 3) physiological stress brought on by high temperatures and oxygen limitation; and 4) physical injury and physiological stress brought on by extreme flows. These have overlapping causes but have all been well documented for Atlantic salmon or closely related species. Our study is unlikely to assign effects to finer categories than these due to the coarse scale of our observations.

2.2.1 time budget, feeding, predation

Atlantic salmon have a well documented seasonal activity pattern with maximum time spent on feeding in the spring, decreasing activity throughout the summer, and a dramatic drop in autumn. Feeding and all other activity is minimal through winter (Cunjak, 1988; Gardiner & Geddes, 1980), increasing again in the early spring. The timing of these seasons varies spatially, with spring feeding beginning prior to ice breakup and ending sooner in northern latitudes. However, this pattern depends on the metabolic status of individuals. Low lipid reserves serve as a trigger for feeding (Finstad, Berg, *et al.*, 2010; Finstad, Ugedal, *et al.*, 2004) and winter feeding carries a survival cost, presumably due to predation. Metabolic requirements related to smolt timing also interact with the time budget. Individuals in their final freshwater winter prior to smolting continue to feed throughout the entire winter season (N. Metcalfe *et al.*, 1988). As allocation of time among feeding and sheltering has been shown to be strongly related to mortality in a closely related system (Biro *et al.*, 2003), I expect this mechanism to apply in our system. Specifically, poor environmental conditions can be expected to reduce growth and result in compensatory feeding either when conditions recover or if lipid reserves are later depleted.

2.2.2   shelter availability

Atlantic salmon are generally territorial and this is also true of shelter locations (Harwood *et al.*, 2002). Attempts by salmon to monopolize shelters result in scarcity of shelters at sufficiently high densities. At the Westbrook, shelter availability is low relative to sites in Norway (*personal communication*) so I expect this mechanism to be important. Measured shelter availability has already been shown to be positively correlated with growth in our system (Davidson *et al.*, 2010) and I expect to see direct and indirect effects on survival through growth due to changes in time budgets. Finally, shelter availability in small freshwater streams is strongly affected by flow. Changes in wet width modify the area of accessible stream bed habitat and can either directly make shelters inaccessible or make them undesirable by disconnecting shelters from feeding areas.

2.2.3   temperature stress

Atlantic salmon, like many salmonids, are cold-water species and their temperature tolerance has been throughly described in the laboratory (J. M. Elliott, 1991). Incipient lethal temperatures, defined as the period of time survivable by 50% of a population, for Atlantic salmon are $22 - 28°$C. The lower end of that range is reached in the peak of summer at our site, suggesting that temperature stress will be a relevant mechanism, but the lack of longer-term laboratory studies on temperature effects makes the strength of the effect uncertain. An additional source of uncertainty is in-stream temperature heterogeneity, which is exploited by fish seeking shelter from extreme temperatures (Breau *et al.*, 2007; Torgersen *et al.*, 1999). Salmon growth is known to be more sensitive to temperature with measured optimum growth temperatures of $17°$C (Bal *et al.*, 2011) for a population from northern France and $16°$C for a pouplation from northwest England. A previous analysis of data from our study site also shows a consistent negative effect of temperature on growth (Davidson *et al.*, 2010).

2.2.4   flow velocity stress

Low flows affect survival through two other mechanisms. First they reduce habitat that reduces shelter and feeding opportunities and can result in increased risk of predation and a deficient energy budget. Second, low flows at times of high temperature stress can further

limit available oxygen and thermal refugia. Direct effects of flow are expected only at high flows when the stream bed is mobilized and holding position becomes energetically expensive. Atlantic salmon are uniquely adapted for holding position in high flows Peake *et al.* (1997) so I do not expect a strong negative effect of flow in our system, particularly for older fish. Young of year (YOY) have more limited swimming ability and Enders *et al.* (2005) showed that high flows do affect feeding behavior, which leads us to expect some stage-specific negative correlation between extreme flows and survival. Accordingly, in another system Jensen & Johnsen (1999) show no relationship between high flow events and survival for over-yearling fish and only a moderate negative effect for young of year.

2.3 Study Description

The Connecticut river watershed is located on the southernmost tip of the known historical range of Atlantic salmon on the North American continent (Parrish *et al.*, 1998). Perhaps salmon populations were always marginal here, which leaves open the question of why use a field study in such a marginal environment. Clearly there are latitudinal gradients in Atlantic salmon biology and if our goal were to parameterize prediction models for northern populations directly on this somewhat extreme environment, we would be hard pressed to justify such an application. Instead, our focus is on transferability and in this case we have much to learn from comparing our general knowledge of Atlantic salmon biology and important parameters to an extreme environment. In addition, some of what makes our Connecticut River system extreme is expected to travel north with climate change, which suggests that our system may provide valuable insight for northern populations. The relevant latitudinal differences are: 1) water temperature; and 2) ice dynamics.

The southern end of the Atlantic salmon range experiences higher summer water temperatures when it is thought to be most likely to affect salmon survival. This difference is muted at the West Brook study site as our site is located in a heavily forested valley, which reduces summer air temperatures and insolation. The study stream also receives stable groundwater inputs in all but the most extreme summers, which tends to increase base flow and provide a buffer against dewatering. Southern environments also generate different ice dynamics. While anchor ice is common in our system, it rarely develops enough to limit habitat availability, which is a concern at northern latitudes. Additionally, while surface ice typically does develop and provides additional shelter in winter, warm winter conditions can

result in limited or non-existent ice cover, which gives predators more access to the stream.

2.3.1 West Brook description

We carried out the mark-recapture fieldwork (42°25 N; 72°39 W) on the West Brook, a 6.3 kilometer long third order tributary of the Connecticut River with a 2% average gradient. Study sections are located 1 kilometer below a reservoir that feeds the stream through an overflow spillway. The study sections are composed of riffle-runs (depth ≤ 45 cm, surface water velocity ≥ 20 cm/s), with 12 identified pools (depth $\geq 45cm$, surface water velocity ≤ 20 cm/s). Each study section is ≈ 20 meters long with a mean wet width of 4.7 meters.

The riparian zone is forested with mixed hardwood trees, which form a dense canopy in the summer. Summer temperatures in this stream are relatively cool (max recorded 21° C, 99th percentile 19.8° C); winter temperatures decline towards the 0-4 degree range in December and begin to rise in early April. Surface and anchor ice formation is common in winter but streamflow is maintained throughout the winter. During summers and in all except four droughts (summers of 1998, 1999, 2001, 2001) the study sections of the stream typically maintain surface flow due to groundwater contributions even when other segments of the stream form disconnected pools. Substrate is primarily cobble and gravel with a mean substrate size of 21.5 cm. Shelter density is low for an Atlantic salmon stream (Davidson *et al.*, 2010).

2.3.2 Stocking techniques

In most years (1997-2000 and 2003) salmon were brought from the White River National Fish Hatchery (US Fish and Wildlife Service, Bethel, Vermont) and stocked into the West Brook as unfed fry (26-28 mm fork length) at a density of 50 fry per 100 meters square in the last two weeks of April or the first week of May. In the years 2001 and 2002, we carried out two rounds of stocking experiments. In 2000, we used two stocking techniques. First, we stocked five families of fish in artificial incubators at ≈ 100 fry per incubator, and allowed to emerge naturally. Second, we stocked at 30 fry per 100 meters square at three time points in the spring (4/13, 4/30, and 5/22) in order to compare the effect of stocking on vital rates. In 2001 we repeated the stocking time experiment with the three time points (4/17, 5/1, 5/17). With the exception of eggs used in incubators, we used unfed fry that had consumed $\approx 92\%$ of their yolk for all stocking. We manipulated water temperature in order to achieve this

17

developmental index despite variation in time of stocking. For stocking using incubators, we fertilized eggs and subsequently kept them in in-stream incubation trays. In total these efforts allowed us to analyze ten overlapping cohorts. Stocking techniques were only varied for other experiments to explore the effect of stocking timing on early fry survival.

2.3.3 Mark-recapture technique

We captured fry and juvenile Atlantic salmon using electrofishing (300-500 volts, unpulsed DC). The portion of the voltage range used depended on stream conditions during sampling. Low flows required a reduced voltage and reduced pursuit time to avoid mortality. High flow periods required a larger voltage to maintain the effectiveness of sampling. We divided the study site into three segments. For 47 central sections, we used block nets (9 mm mesh) to contain fish in each $\approx$ 20 m section. Electrofishing always started at the lowest sections and moved upstream. Typically, one person carried an electrofishing backpack and a net, while two others carried only nets. To capture fish, the person carrying the electrofishing backpack, inserted the probe into the water, turned it on, and swept it through a section of the stream. Fish within a few feet of the probe became unable to swim and we captured and removed them to water-filled buckets. The team of three started at the downstream end of each section and worked its way across and then up the stream. The end of the pass occurred when the team arrived at the upstream end; at each sample we completed two passes for each of the central 47 sections. After each pass we checked the block nets for stunned or trapped fish. We used three sets of nets and leap-frogged the block nets upstream to minimize the possibility that a fish might move down between electrofishing passes. For the peripheral sections upstream and downstream of the central segment we started at the downstream end and moved upstream, but we only completed one pass per section and we did not use block nets.

All captured Atlantic salmon over 55 mm were tagged intraperitoneally by making a small incision between the pectoral fins and inserting a passive integrated transponder (PIT) tag. Scales were taken from all untagged fish to determine age. This technique resulted in high retention ($\geq$ 99%) and low mortality ($\leq$ 5%) (Gries & B. H. Letcher, 2002). Any Atlantic salmon captured were scanned for PIT tags, and their length (+/-1mm) and weight were recorded using field computers (Juniper Systems, Logan, Utah). Recaptures were used to establish histories of when individuals were known to have survived as well as to estimate

recapture probabilities, which provides information on survival after last recapture.

Sampling by electrofishing rarely results in some fish mortality. Dead fish are found either during sampling or when block nets are checked for fish prior to removal. The fate (death) of these fish is recorded so that it is available for modeling.

2.3.4   Smolt trap techniques

From 1998 to 2001 we used a picket weir (Anderson & MacDonald, 1978) to trap emigrating smolts. We installed the trap in April once temperatures reached 5°C and dismantled it two weeks after the last smolt capture. The trap was cleaned and checked for fish twice daily. The trap was effective in 1999 and 2001 but did not operate effectively in 1998 and 2000 due to high flows. From 2001 to 2003 a combined screw trap (Scace *et al.*, 2007) and picket weir were used to trap emigrating smolts. The smolt traps provided limited information on smolt timing and probability, as well as identifying some individuals that left the study site.

2.3.5   Antenna system

Emigration of tagged fish from the study site was monitored using a stationary PIT tag antenna system (Castro-Santos *et al.*, 1996). The configuration of antennas at the West Brook changed over the years but generally included a pair of antennas at the downstream end of the study site and one at the upstream end. Fish last observed at an antenna and never again were considered permanent emigrants and their fate (survival or death) was not modeled past that time point. From May 2001 to February 2008, a pair of antennas (1.1 m x 0.4 m) were operating at the downstream end of the study site (Zydlewski *et al.*, 2006). A single antenna (1.1 m x 0.4 m) operated at the top of the study site. The antenna system was operated in order to disentangle true survival from emigration and remove the typical ambiguity between environmental effects on emigration and environmental effects on survival. Individuals that were last seen at an antenna were censored and their life histories were not modeled past that time point.

2.3.6   Discrete season boundaries

Over a decade of quarterly mark-recapture sampling, electrofishing samples generally started in mid-March, mid June, early August, and late November. These sampling dates correspond

Table 1: Definitions of biological seasons for Atlantic salmon.

season	start	stop
spring	68	140
summer	141	247
autumn	248	335
winter	336	67

to important time-points for the interaction of Atlantic salmon and their environment. The mid-March sample is able to recapture overwintering salmon before they are able to smolt. The June sample captures most non-smolts, early August is a time with increased mobility by mature salmon, and the late November sample captures individuals immediately before they become unavailable due to stream conditions (snow, ice and decreased response to electrofishing). For practical reasons the actual timing of samples varied, making it difficult to assign samples to seasons. Therefore I calculated season boundaries as follows:

1. Use k-means clustering (4 clusters) to assign samples four groups based on their start dates.

2. Label each sample according to its calculated season (group).

3. Visually inspect the distribution of sample start dates grouped by season to assess the number of early outliers.

4. Choose the start date of the third earliest sample within each season to represent the season start date.

By following this procedure I assigned Julian days to seasons as shown in table 1:

This classification of seasons is faithful to the seasonal nature of the recapture design but insensitive to the vagaries of fieldwork. For some models I also used the day of the year directly as a covariate indicative of seasonality.

2.3.7 Deviation from the typical season

To characterize environmental variation at the study site, we used a water temperature probe and data logger at the downstream end of the study site. We also used a water depth logger and stage-discharge relationship to calculate an index of stream discharge (m^3s^{-1}). These covariates were measured at a high rate and summarized to daily means. Discharge was highly skewed towards large values so I used a $\log_{10}$ transform to generate a convenient covariate. Both covariates include some season-specific patterns with 1) temperature near freezing for part of the winter, increasing in the spring and through the summer, and again falling in autumn and early winter; and 2) flows generally increasing in the transition from winter to early spring, dropping in through spring and summer, and typically rising again in early autumn and winter. My goal was to characterize environmental variation around these typical patterns, which I did as follows:

1. Calculate the Julian day for each measurement time-point from all years of the study.

2. Fit a loess smoother with Julian day as the predictor and the daily average temperature or log discharge as the dependent.

3. Adjust the smoothing parameter so that only seasonal trends were captured by the smoother.

4. Predict a mean daily temperature or log discharge for each Julian day.

5. Subtract the predicted value from the daily average temperature or log discharge.

The resulting covariate has a value of zero (in its original units) when the average value for that day is at the expected value for that time of year. Positive values indicate a warmer than typical day, or higher than typical discharge for a given part of the year. To use these covariate values for our seasonal modeling, I calculated the mean of these values between individual recaptures or recapture occasions, these seasonal values became our flow (F) and temperature (T) covariates. For the recapture model only, I also used the flow on the day of capture (F_d) calculated as above.

2.3.8 Environmental covariate patterns

High rate covariates present a special challenge in modeling environmental responses with coarse data. We observe survival—but never death—four times a year at most, however, we observe environmental data at least hourly. Matching these two scales means deciding *a priori* which summaries to use. There is support in the literature for a variety of summaries, many more than we could possibly relate to our data.

In order to simplify this part of the modeling task I chose to use a seasonal component as well as average deviation from the seasonal component as defined below. Variance in environmental covariates can be equally important in general (Lawson *et al.*, 2015). In our system variance can also be important as there are documented responses to short-term high-flow and high-temperature events. The techniques I develop could be applied to other summary statistics but they are beyond the focus of this dissertation.

Both flow and temperature showed seasonal variation (figure 1), with the fastest changes in temperature occuring during spring and autumn. Winter temperatures were consistently in the 0-4 degree range, spring showed a rapid rise ($\approx 0.1°C$ per day), which continued into the summer. Temperature typically peaked before the final third of the summer season and dropped during autumn ($\approx 0.1°C$ per day). Winter flows were typically consistent around ≈ 0.3 m^3/s, flow rose quickly and peaked mid-spring, primarily due to snow melt and limited absorption capacity due to frozen soils. Peak spring flows were around ≈ 0.6 m^3/s. Flows decreased over the course of the summer with an initial period similar to winter and a final low of ≈ 0.03 m^3/s at the end of the season. In autumn, flows typically returned slowly to winter levels. Environmental covariates calculated as deviations from these seasonal patterns indicate whether a particular season was *exceptionally* wet, dry, hot or cold. This means that in summer a high discharge covariate might indicate more favorable conditions whereas in spring a high discharge covariate might indicate flood conditions with potentially negative impacts.

As designed, seasonal boundaries clearly divide major portions of the yearly environmental cycle. Maximum observed summer temperatures are clearly below the known thermal maxima of Atlantic salmon (J. M. Elliott & J. A. Elliott, 2010) but above the optimum growth temperature (15°C, J. M. Elliott & Hurley (1997)). Winter temperatures clearly dipped into the no-growth zone (J. M. Elliott & Hurley, 1997).
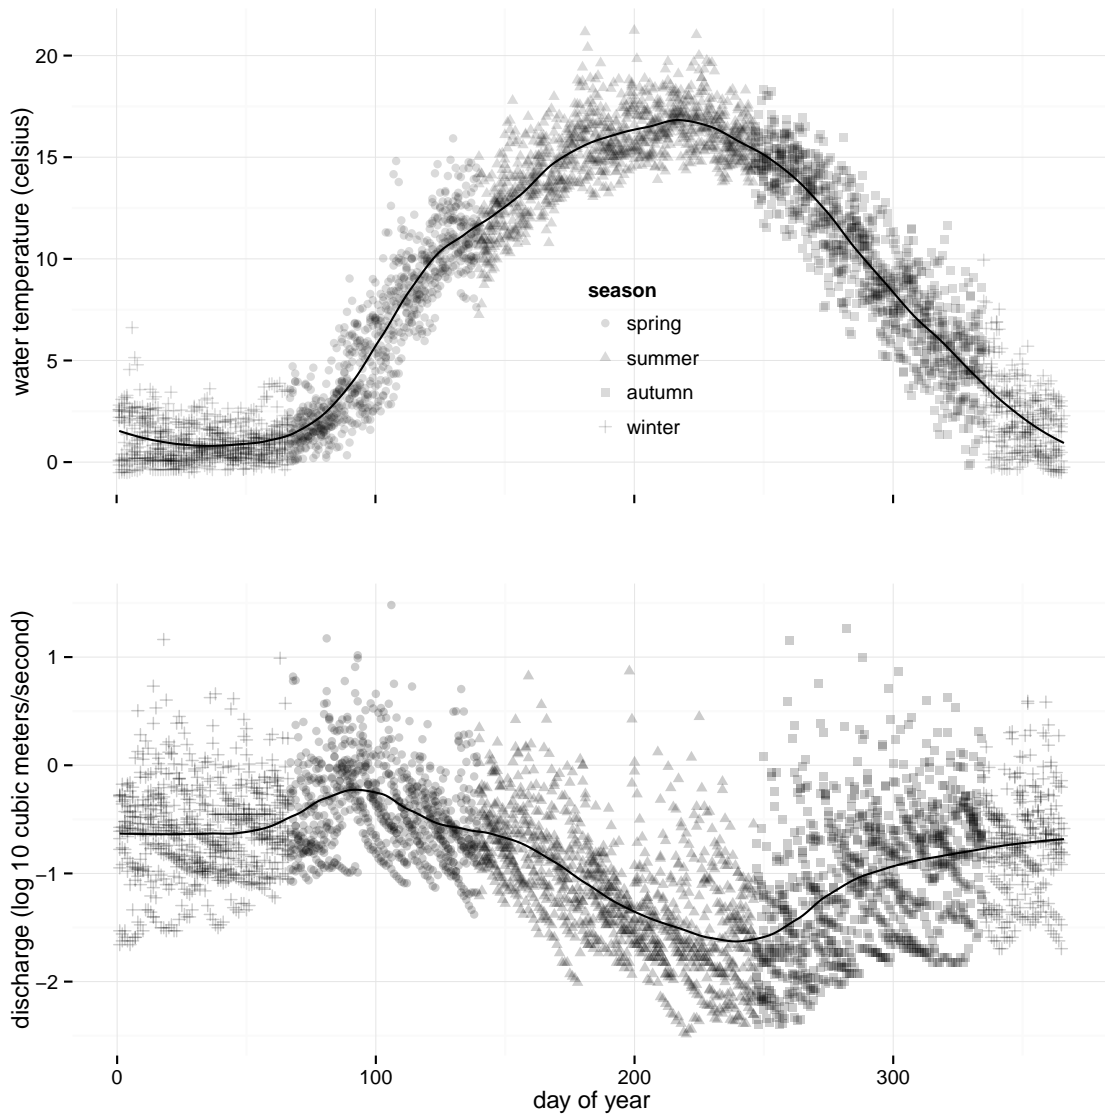
Figure 1: West Brook water temperature and flow data show a strong seasonal pattern. Minimum winter temperatures are typically in the no growth range for Atlantic salmon. Maximum summer temperatures are below any lethal levels from the literature. Discharge also shows a strong seasonal pattern with high flows in spring and decreasing flows throughout the summer. Flows return to average levels over the winter until the spring flush.

23

CHAPTER 3

MARK-RECAPTURE IN CONTINUOUS TIME

To use mark-recapture data to estimate vital rates and their sensitivity to environmental effects, ecologists commonly rely on the Cormack-Jolly-Seber (CJS) model (Cormack, 1964; Darroch, 1958, 1959; Jolly, 1965; Seber, 1965) and its extensions. Most of the development in the basic model is summarized in Lebreton *et al.* (1992) with Barker (1997) and Kendall, Nichols & J. E. Hines (1997) introducing important extensions. Barker (1997) introduced a more holistic view of mark-recapture data by incorporating (live) recapture, (live) resighting, and (dead) recovery data and was an early forerunner of current work on integrating varied data types (Schaub, Gimenez, *et al.*, 2007). The original CJS model was restricted to consider only two states—alive and dead—and Kendall, Nichols & J. E. Hines (1997) extended this to multiple, sometimes unobserved, states. These models are supported by software with Program MARK and various BUGS-language implementations being most common. For a number of reasons we were not able to use these options (see appendix A for details) and relied on Stan (Stan Development Team, 2014) instead.

The model I apply is best thought of as a multi-state extension of the CJS models described in Lebreton *et al.* (1992). Our primary data source is the mark-recapture survey, described fully in section 2.3. The model ignores individuals prior to their first capture. When an individual is tagged it enters the population and can only exit the population when it 1) dies, with probability $1 - \phi$; 2) emigrates and is detected by antenna, with probability ρ; or 3) emigrates and also smolts, with probability $\rho \times \alpha$. The smolt trap and emigration antennas were both assumed to be 100% effective up until the last successful recapture, but we included undetected emigration and/or smolting as possible outcomes after the final successful recapture. Because data on smolting is limited and emigration is detected directly, I used a simple seasonal model for these components. I assumed both were absorbing states. Smolting is certainly a well identified absorbing state, whereas emigration is defined as the last re-sighting at a downstream antenna, which makes emigration an absorbing state by definition. I do not include a model for temporary emigration.

In addition to possible detection at emigration and smolting, the CJS model assumes that recaptures of individuals are also attempted at specific time points, labeled "occasions". I define a random variable X to describe whether a successful capture occurs ($X = 1$) or

not ($X = 0$). At each occasion, j, an individual, i, has a well-defined recapture probability $p_{i,j}$ of recapture. When an individual is known to be available (e.g., is captured later) the model for recapture is: $X \sim \mathrm{Bern}(p_{i,j})$. When an individual was last detected alive, and is not subsequently detected as a smolt or an emigrant, the remaining possibilities are that it survived and was not detected; or it died. With each *attempted* but failed observations it becomes more likely that the true outcome was death.

The CJS model explicitly represents each process and allowed me to separate parameters for the survival process we are interested from parameters of the (unimportant) recapture process. Separating component models also made it possible to elaborate the survival model and measure environmental effects more directly rather than relying on observed patterns of recaptures. The main cost to using this model is some ambiguity about whether a given data set provides enough information to resolve the confounding between survival, recapture, emigration, and smolting parameters. I address this ambiguity in the results section along with the substantive survival estimates.

3.1 Methods

3.1.1 Notational interlude

This chapter and the chapters that follow deal primarily with events, the timing of events, states, and state changes for well-defined units. I define the notation used to refer to these entities here and use this notation throughout the rest of the dissertation. As much as possible I follow conventions from the statistical literature, though I try to avoid the more ambiguous conventions.

A unit might be an individual in a mark-recapture study that enters the study at a specific time and remains alive until death or some other event occurs. The unit enters the study in a given state, and events that affect the unit are observed. Individual units are indexed by i and the time of an observation is indexed by j.

I indicate random variables by using capital letters, typically X, Y, and Z. Random variables are associated with distributions that describe how their outcomes arise, for example:

$$X \sim \mathrm{Bernoulli}(p)$$

When I discuss a series of similar distributions and random variables, they are distinguished by indexing the random variables and the associated parameters, as in:

$$X_1 \sim \text{Bern}(p_1)$$
$$X_2 \sim \text{Bern}(p_2)$$
$$\ldots$$
$$X_M \sim \text{Bern}(p_M)$$

Realizations of random variables for a given unit and observation are denoted by the corresponding letter in lower case: $x_{i,j}$ or $y_{i,j}$. State changes (transitions) are statements about the transition from one state to another given the probability model and parameters. Parameters are either indicated explicitly or elided when obvious given the context and replaced by an ellipsis.

$$Pr[X = x_{i,j}|\theta] = Pr[X = x_{i,j}|\ldots]$$

In the present context where we follow units as in a longitudinal study, we deal with the state of a unit at multiple time points. The initial states of the units are typically observed when they enter the study and they remain in that state until the time of a state change. These are typically written as conditional probabilities, for example:

$$Pr[X = x_{i,j}|X = x_{i,j-1}, \theta] = Pr[X = x_{i,j}|X = x_{i,j-1}, \ldots]$$

When only one unit is discussed, the index i is dropped. Each observation on a single unit x_j is associated with a time point also indexed by j, written as either t_j or s_j. Each observation after the first is also associated with a duration, written d_j, so that the duration from t_0 to t_1 is written d_0.

In likelihood or probability density calculations it is often useful to distinguish between variables that will always be fixed in the local context and those that may be integrated out as nuisance parameters, estimated, or otherwise the focus of manipulations in a given equation. For variables that we fix to a single value we use s, and for focal variables we use t. When there is no useful distinction to be made we use t only. When only two time points are in play we use s for the first and t for the second.

26

In all cases indexing of units takes values from 1 to N ($\{i \in \mathbf{N} | i \leq N\}$), and indexing of observations runs from 1 to J ($\{i \in \mathbf{N} | i \leq J\}$). I follow this pattern for any other indexes so that a group index k would run from 1 to K and so on.

3.1.2 Probability model

The Cormack Jolly Seber (CJS) model can be written based on the state Z of each individual, i, at time points $1, \ldots, J$. Z is defined as 1 when alive and available for capture and as 0 otherwise. The model and can be written as:

$$Z_{i,j} | Z_{i,j-1} \sim \text{Bern}(Z_{i,j-1} \times \phi_{i,j-1} \times (1 - \alpha_{i,j-1}) \times (1 - \rho_{i,j-1})) \tag{1}$$

With parameters defined as per-individual and per-occasion survival ($\phi_{i,t}$), recapture ($p_{i,t}$), emigration ($\rho_{i,t}$), and smolting ($\alpha_{i,t}$) probabilities. To complete this I add a simple model for recapture outcome X, which is defined as 1 when a recapture is successful and 0 otherwise. This outcome is conditional on the state as only alive resident individuals are available for recapture:

$$X_{i,j} | Z_{i,j} \sim \text{Bern}(Z_{i,j} \times p_{i,j}) \tag{2}$$

This large number of parameters is not identifiable so specific models relate these parameters to fewer higher-level parameters. The definitions of these first level parameters are:

ϕ: Survival over a standard quarter-year.

p: Probability of recapture from electrofishing at a point in time.

ρ: Probability of emigration at a point in time.

α: Probability of smolting at a point in time.

The survival parameter is defined for a period of time because the data collected—recaptures—directly informs us about the event: survival. Recapture is not dependent on the duration of time since the last event so it is associated with an instant. It may be more correct to think of the probability of smolting or, especially, emigration as something that depends on the duration of time considered but the data we collect does not inform us about

when the decision to smolt or emigrate is made. Therefore we only associate the probability of smolting with conditions at a point in time. I use these parameters indexed by event (i: individual, t: time) to construct per-event log likelihoods as described in table 2, which are then summed to construct the model likelihood.

There is no closed form for the likelihood of events after the final observation when the state is unknown and the only information brought by these terms is in terms of the probability of *not* observing the individual within the study site *if* it happened to survive. The recursive relationship for the probability of a non-observation conditional on survival is:

$$Pr[X_t = 0] = Pr[X_j = 0 | Z_j = 1] + Pr[Z_j = 0 | Z_{j-1} = 1] \tag{3}$$

This recursive relationship is used to sum over all the possible intervals that the individual could have survived in order to calculate the likelihood contribution and each sum is terminated by a death term $(1 - \phi)$. In practice the terms decrease in magnitude quickly if p is relatively high, making it practical to carry the sum out over a small number of intervals after the last observation. An efficient non-recursive version of the sum is implemented in the Stan program associated with this dissertation (appendix B). Note that as long as for each individual the events are contiguous and well-defined, this model can be reinterpreted as a continuous-time survival model.

3.1.3 Data classification

While data and events (e.g., recapture or emigration) were classified based on their associate timepoint, processes such as survival occur over an interval of time. To match this modeling framework we calculated the interval corresponding to each classified state and calculated further covariates (e.g., season or average temperature) based on this interval. For a pair of states, $X_i(t^*)$ and $X_i(t)$, we defined the relevant interval m, as $\{m : t^* < m \leq t\}$. The season assigned to an observation was based on this interval and due to events at season boundaries each interval had an unambiguous season designation. The average environmental covariates were also calculated for these intervals. Finally, the duration, d_j, of the interval was calculated as $t - t*$ and used to scale survival probabilities to survival over a standard quarter-year.

Our data on survival consisted of recaptures of live individuals, some known deaths, observed emigration, and smolt captures. The recapture process was imperfect, leading us to use the Cormack-Jolly-Seber model (Cormack, 1964; Darroch, 1958, 1959; Jolly, 1965; Seber, 1965) with modifications to account for both emigration and smolting. The state (X) of each tagged individual (i) at each time point (j), is classified as follows:

$$X_{i,j} = \begin{cases} 1, & \text{if individual } i \text{ is recaptured alive at timepoint } j. \\ 2, & \text{if individual } i \text{ is uncaptured but known alive at timepoint } j. \\ 3, & \text{if individual } i \text{ is last detected at a boundary at timepoint } j. \\ 4, & \text{if individual } i \text{ is last detected at a smolt trap at timepoint } j. \\ 5, & \text{if individual } i \text{ is known alive at timepoint } j. \\ 6, & \text{if individual } i \text{ is in an unknown state at timepoint } j. \\ 7, & \text{if individual } i \text{ is known censored (previously dead/emigrated/etc...) at timepoint } j. \end{cases}$$

Based on the data collected and the existence of an 'unknown' state, as in the classical CJS model, the individual state can be classified at any time point. This classification is applied for all time points after tagging until the individual's record is censored (i.e.-observed death, smolting, or emigration) or until one and a half years after final capture. The choice of time points is in part dictated by data since the time of capture, censoring, emigration, or smolting are observed. For the remaining states, time points were assigned as follows:

$X = 2$: Uncaptured individuals were assigned a timepoint uniformly distributed between the first and last capture of a given capture occasion. This state and timepoint was only generated if an individual was *not* captured during a given occasion.

$X = 5$: Regardless of any other captures, a timepoint was generated for each individual at each season boundary. This made it possible to unambiguously assign each interval to a particular discrete season by dividing intervals that would otherwise cross a season boundary.

$X = 6$: When individual state was unknown, timepoints were still generated to allow for the recursive CJS likelihood calculation. Since model parameters depend on season and age classifications, these timepoints were also generated at season boundaries.

This classification defines all the *events* that we model and they are described in table 2 with their per-event likelihoods.

Table 2: Mark-recapture per-event likelihoods. The event class matches the data classification in the text. All parameters are individual and time-specific but these subscripts are dropped in the table for clarity. The parameters are ϕ for survival, p for recapture, α for smolting, and ρ for emigration. For individuals in an unknown state we follow the usual solution and calculate the likelihood using a recursive relationship (see equation 3).

event class	likelihood	description
1	$\phi \times (1 - \rho) \times (1 - \alpha) \times p$	individual is recaptured alive within the site
2	$\phi \times (1 - \rho) \times (1 - \alpha) \times (1 - p)$	individual is know alive but *not* recaptured
3	$\phi \times \rho$	individual is known alive and observed to emigrate
4	$\phi \times \rho \times \alpha$	individual is know alive and smolted, emigration is implied.
5	$\phi \times (1 - \rho) \times (1 - \alpha)$	individual is known alive and still within the site
6	$f(\phi, p, \alpha, \rho)$	individual state is unknown, see equation 3 for likelihood.
7	0	individual is censored (known smolted, emigrated, or dead).

Table 3: Model structure specification for survival (ϕ), recapture (p), smolting (α), and emigration (ρ) parameters. The notation is as follows: 1) subscripts indicate a discrete effect; 2) terms in parentheses indicate a continuous relationship; 3) terms such as $A \times B$ indicate both individual factor effects and interactions; 4) terms such as $A : B$ indicate interactions only; 5) terms such as A, B, K, with K constant, indicate a multi-dimensional spline; and 6) terms such as A, K, with K constant, indicate a single-dimensional spline with K knots. Each component model is transformed through the inverse logistic transformation to generate an event probability.

Model#	Survival	Recapture	Smolting	Emigration
1	ϕ_S	$p(d,4) + p(dT) + p(d\log(F))$	α_A	ρ_S
2	$\phi_{A \times S}$	$p(d,4) + p(dT) + p(d\log(F))$	α_A	ρ_S
3	$\phi_{A \times S}$	$p(d,4) + p(dT) + p(d\log(F)) + p_{OCC}$	α_A	ρ_S
4	$\phi_{A \times S} + \phi_{A+A:S}(T) + \phi_{A+A:S}(F) + \phi_S(T \times F)$	$p(d,4) + p(dT) + p(d\log(F))$	α_A	ρ_S
5	$\phi_{A \times S} + \phi_{A+A:S}(T) + \phi_{A+A:S}(F) + \phi_S(T \times F)$	$p(d,4) + p(dT) + p(d\log(F)) + p_{OCC}$	α_A	ρ_S
6	$\phi_{A \times S} + \phi_{A+A:S}(T)^2 + \phi_{A+A:S}(F)^2 + \phi_S(T \times F)^2$	$p(d,4) + p(dT) + p(d\log(F)) + p_{OCC}$	α_A	ρ_S

3.2 Model structure

Due to the computational requirements of complex models with a large data set and the broad knowledge base about mark-recapture techniques and Atlantic salmon biology, I chose a set of *a priori* models. I estimated the models sequentially, generally moving from simpler to more complex structures in order to discover computational problems individually and to generate starting values for more complex models. In all cases I used weak priors that encompassed the plausible set of effect sizes. The models are summarized in table 3. Each component model is a linear model for an event probability and I chose to use the inverse logistic transformation for all components to move from the linear model coefficients to the event probability unscaled event probability.

Despite their complexity, these models are basic models for known components of our system. Even the most complex model could be modified to capture further effects that are known to exist from laboratory studies. I limit the analysis to these models as they can be used to address the questions at hand and I expect the data to contain sufficient information to estimate these effects. The simple smolting and emigration models allow us to estimate the overall importance of these sources of apparent mortality but in the case of smolting very limited data are available. In the case of emigration most departures are observed and departing individuals are censored, which limits the effect of model mis-specification on survival and recapture estimates. Model components are described below, with the following notation:

1. subscripts and square brackets indicate a discrete effect;

2. terms in parentheses indicate a continuous relationship;

3. terms such as $A \times B$ indicate both individual factor effects and interactions;

4. terms such as $A : B$ indicate interactions only;

5. terms such as A, K, with K constant, indicate a one-dimensional spline with K knots; and

6. terms such as A, B, K, with K constant, indicate a multi-dimensional spline with K knots per dimension.

3.2.1 Recapture

The goal of the recapture model is to 1) account for missing observations and 2) to develop a parametric structure for this accounting that can be used to evaluate the coherence of the reporting model vis-à-vis field experience. The relevant field experience for the West Brook system is as follows:

1. In the cool seasons, particularly winter, salmon activity is diminished and correlated with temperature. Extremely low temperatures result in fish sheltering in the substrate and a slow response to electroshocking, which results in lower recapture rates.

2. In the summer season, low flows result in fish sheltering in the substrate and restrict the permissible pursuit of individuals that initially respond to electrofishing due to an increased risk of mortality. This results in lower recapture rates for low-flow seasons in the summer.

3. In the high-flow spring season in particular, but also in any other high-flow situation, the increased volume of the particular stream sections makes it difficult to capture individuals that are sighted when they initially respond to electrofishing. Decreased recapture rates are expected in high absolute flow setting.

4. Environmental covariates aside, seasonal variation in activity and response to electrofishing affects recapture rates.

5. Uncharacterized environmental variation as well as the experience of the field crews affect recapture rates.

To capture these effects, I constructed a set of environmental, seasonal, and occasion-specific effects on recapture rates.

It is common for mark-recapture studies to capture seasonal variation in recapture probability using a season indicator. The present study was carried out over the course of a decade in a wide range of environmental conditions. Environmental conditions resulted in important variation in sample start and end dates within seasons. To avoid classifying recaptures from a single occasion into multiple "seasons", I used a seasonal radial spline on the Julian day of recapture, d, to model recapture probability. The radial spline uses a

parametric basis function, f, to define the intercept β_k around a given time of year, mu_k, and a smooth function in between these points.

$$p(d, K) : \text{logit}(p_{i,t}) = \sum_{k=1}^{K} \beta_k \times f(d_{i,t}, \mu_k, \sigma^2) \tag{4}$$

In addition to seasonal variation, I expect temperature effects on recapture as fish respond to cold water temperatures by becoming lethargic, leading us to use centered unstandardized temperature as a covariate for recapture. Flow on the day of the recapture has a direct impact on recapture probability so I used the centered log average flow for that day as the covariate.

$$p(dT) + p(d\log(F)) : \text{logit}(p_{i,t}) = \ldots + \beta_{dT} \times dT_{i,t} + \beta_{d\log(F)} \times d\log(F)_{i,t} \tag{5}$$

Many other factors, including the number of sampling team members, the harshness of conditions, and the experience of available team members affect recapture probability. These effects can not be estimate independently for all occasions due to identifiability issues but we were able to fit a per-occasion random effect (p_{OCC}) to describe this variation.

$$p_{OCC} : \text{logit}(p_{i,t}) = \ldots + \beta_{OCC[i,t]} \tag{6}$$

These occasion-specific effects were constrained to be Cauchy-distributed with mean zero and estimated variance. The heavy-tailed distribution allowed recapture probability to be estimated as near-zero for failed samples when stream conditions were very poor. For more common recapture conditions this distribution applies some shrinkage to the estimated recapture effects. Relying on shrinkage is appropriate for this study due to the limited information available for estimating recapture probability in non-robust mark-recapture study designs.

3.2.2 Survival

The goal of the survival model is to characterize survival probability with respect to age structure, seasonality, and especially environmental components. Environmental conditions

change over the course of the year in a predictable way (figure 1) and are known to combine with seasonal changes in behavior and physiology in ways that affect survival. I used the standard approach of assigning each interval to a specific season and estimating a season-specific intercept for survival across all age classes.

$$
\begin{aligned}
\phi_S : \text{logit}(\phi_{i,t}) =& \beta_{s[i,t]} \\
\phi_{A \times S} : \text{logit}(\phi_{i,t}) =& \beta_{s[i,t]} + \beta_{a_{0+}} \times (A_{i,t} = 0^+) + \\
& \beta_{a_{2+}} \times (A_{i,t} = 2^+) + \beta_{a_{3+}} \times (A_{i,t} = 3^+) + \ldots + \\
& \beta_{a_{J+}} \times (A_{i,t} = J^+) + \\
& \beta_{x,s \neq 2, a \neq 1+}
\end{aligned}
$$

Age 1^+ is the intercept in this scheme, with an offset for age 0^+ and a separate offset for each of the older ages. Initially I used a series of sequentially applied offsets for older ages to maximize the amount of information contributing to the younger age classes. In early models age effects were clearly not very important and I simplified the age structure.

Due to the seasonal changes in average environmental conditions, I expected the response to environmental conditions to also vary seasonally. I structured our environmental models following the same age and season structure as the intercept models for survival. I present the equations used to calculate coefficients here for completeness.

$$\phi_{A+A:S}(T) : \mathrm{logit}(\phi_{i,t}) = \ldots +$$

$$\beta_{a_{0+}} \times (A_{i,t} = 0^+) \times T + \beta_{a_{2+}} \times (A_{i,t} = 2^+) \times T +$$

$$\beta_{a_{3+}} \times (A_{i,t} = 3^+) \times T + \ldots + \beta_{a_{J+}} \times (A_{i,t} = J^+) \times T +$$

$$\beta_{x,s \neq 2, a \neq 1^+} \times T$$

$$\phi_{A+A:S}(F) : \mathrm{logit}(\phi_{i,t}) = \ldots +$$

$$\beta_{a_{0+}} \times (A_{i,t} = 0^+) \times F + \beta_{a_{2+}} \times (A_{i,t} = 2^+) \times F +$$

$$\beta_{a_{3+}} \times (A_{i,t} = 3^+) \times F + \ldots + \beta_{a_{J+}} \times (A_{i,t} = J^+) \times F +$$

$$\beta_{x,s \neq 2, a \neq 1^+} \times F$$

$$\phi_S(T \times F) : \mathrm{logit}(\phi_{i,t}) = \ldots + \beta_s \times T \times F$$

$$\phi_S(T, F, 4) : \mathrm{logit}(\phi_{i,t}) = \ldots + \sum_{k=1}^{4} \sum_{m=1}^{4} \beta_{k,m} N((T, F), \mu_{k,m}, \sigma^2)$$

The structure of the age effects and interactions mirrors that for the intercepts. The only addition was a seasonal interaction term between flow and temperature.

3.2.3 Smolting and emigration

Data on smolting is very limited relative to the rest of the study. I used a seasonal and an age-specific component to account for the proportion of individuals lost from the study due to smolting. Smolting probability was set to zero for all non-spring seasons. Smolting was estimated independently for age 2+ and age 3+ and older fish. No 1+ smolts were observed at our site and few fish older than age 3+ were observed, which led us to aggregate all fish age 3 and older. Data on emigration are more complete but do not span the entire study duration which led me to use a simple seasonal model to account for emigration probability based on antenna detections.

3.3 Mark-recapture global model results

3.3.1 Recapture estimates

For the relevant times of the year, recapture estimates are high enough ($p \gtrsim 0.5$) to expect that a mark-recapture model will be effective for estimating survival parameters. When

sampling occasions are by discrete seasons, recapture estimates are consistent regardless of model structure. The simple 4-knot spline recapture model also generally reproduces the recapture discrete seasonal recapture pattern (results not shown). The addition of occasion-specific Cauchy-distributed random effects generates the same pattern but with an exaggerated peak in summer recapture probability (figure 2). The presence of clear outliers in recapture probability (figure 3) results from samples which were aborted when conditions were too likely to result in fish mortality.

Although the occasion-level effects are well-estimated for the core study occasions when sample sizes are adequate, they do display a seasonal pattern with negative winter effects and positive summer/autumn effects. This suggests that while the combined recapture probability may be approximately accurately estimated, a 4-knot spline may not be flexible enough to capture seasonality of recapture and testing 5 or 6 knot splines may prove an effective method for maintaining the conceptual separation between seasonality and occasion-specific effects. Based on a variety of recapture models, including a simple seasonal model, as well as a 4 knot and pathological 8 knot model, I did not see any effect on estimates or survival parameters from the choice of the recapture model structure (results not shown). The lack of sensitivity suggests that with relatively high recapture probabilities, the recapture histories themselves provide sufficient information on when individuals died for any given occasion—likely within the season or two after their last capture.

In addition to solid estimates of per-occasion recapture probability, our model also produced estimates of environmental effects on recapture. Flow had a negative effect on recapture (figure 2) and temperature had only a weak effect on recapture.

3.3.2   Seasonal survival estimates

Regardless of the model employed, seasonal survival is relatively high (≈ 0.75) and strong seasonal patterns are apparent. Survival is highest in summer and winter and lower survival in spring and autumn regardless of age category (figure 4). This seasonal pattern is very robust with regards to model choice and *all* age-by-season models are shown in figure 4 with varying shading to emphasize this point. In contrast to the strong seasonal pattern, the full panel of models shows very limited variation in average survival by age. These seasonal intercepts combine with season-specific environmental effects to determine overall survival under specific environmental condition (figure 6)

Environmental effects were estimated both with and without interactions using a variety of recapture models. Because interactions were ultimately important, I do not present our results about environmental sensitivity from models without interactions. Interaction terms are generally distinct from zero, well enough estimated that MCMC sampling variation does not affect qualitative survival patterns. Two dimensional plots of survival as a function of temperature and flow result in clear patterns of interaction (figure 5), with temperature having a positive effect on survival in all seasons except for summer at low flows. At high flows, temperature has a negative effect on survival in winter and spring, temperature is unimportant in autumn, and it has a strong positive effect in summer.

With complex interactions the plots of individual effects become more difficult to interpret and it can be helpful to look at a response surface instead (figure 6). In spring the surface shows little response of survival to temperature at average flows, but a negative effect of temperature at high flows and a positive effect of temperature at low flows. In summer this pattern reverses completely. Typical summer conditions still show only a weak effect of temperature but high flow conditions result in a positive effect of temperature on survival and low flow conditions result in a negative effect of temperature on survival. In autumn, temperature has a weakly positive effect at average and high flow conditions, but increasing temperatures have a strongly positive effect at low flow conditions. Unlike the other seasons, average autumn are not at the saddle point of the surface. The winter surface mirrors the spring with weak temperature effects at average flows, a negative temperature effect at high flows, and a positive temperature effect at low flows. This complex set of surfaces is well supported by the data. Much of the covariate space where effects are estimated contains multiple resightings of individuals (figure 7). Results for the final model are presented here (figures 6 and 7), with results for simpler models in the appendix. None of the simpler models presented in the appendix gave results that contradict the final model.

The basic survival surface patterns described above occur in the context of seasonal changes which are masked by z-score transformations of the flow and temperature variables. A hot dry spring season refers to different environmental conditions than a hot dry summer. The following presentation of results introduces these seasonal reference points.

Spring flows generally increase above winter levels, driven by snowmelt combined with rainfall before returning to winter levels. Average spring survival is relatively low (≈ 0.72). Environmental effects in spring suggest that the best conditions occur when high temperatures combine with low flows (an early spring following a low-snowpack winter) or low

temperatures with high flows (a slow spring with a decreased rate of temperature increase, conditions that would typically still result in a melting snow pack and disappearing river ice). High spring temperatures combined with high flows are relatively rare (figure 7) and it is difficult to evaluate the effect of such conditions. However, low temperature dry spring seasons are relatively common and result in strongly decreased survival. These conditions would imply an extended winter season and indeed the overall pattern of spring effects matches the pattern of winter effects.

Summer is a season of decreasing flows and increasing/peaking temperatures with relatively high average survival ($\approx .84$). Relatively low summer temperatures combined with low flows (cool dry summer) as well as higher temperatures combined with high flows (warm wet summer) both result in good survival. Hot dry summers as well as cool wet summers result in strongly diminished survival.

Autumn is a season of increasing flows and decreasing temperatures with relatively low ($\approx .68$) average survival. Environmental effects in autumn are weakest but still important. Survival is highest in a warm wet autumn and decreases with drier conditions such as would be expected following a dry summer or when the onset of higher precipitation is delayed. Cool autumn conditions, which would generally be brought on by an early winter, also decrease survival.

Winter flows and temperatures are stable and average survival ($\approx .8$) is relatively high. Like spring, cold and dry winters or warm and wet winters both result in decreased survival, while cold wet conditions and warm dry conditions improve survival. Winter environmental variation tends towards the kind that is bad for survival (cold/dry or warm/wet) and therefore winter seasons that depart from the typical are unlikely to improve survival.

In order to compare these results to the more flexible local regression model I defer discussion of these substantive results to the final discussion chapter.
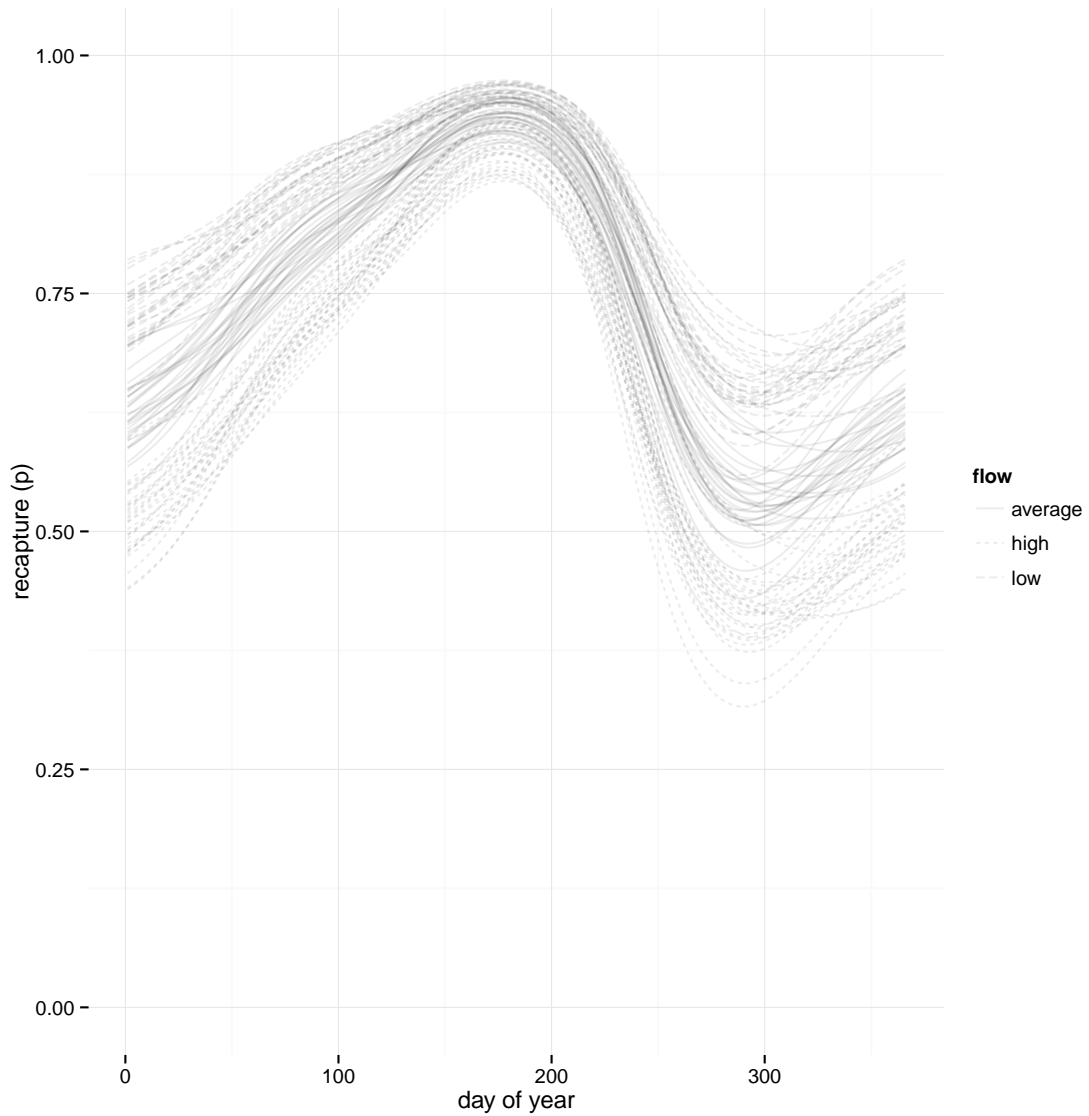
Figure 2: Seasonal estimates of recapture probability. Results presented here are from the final model. Each line represents one draw from the posterior distribution of the effect with low stream flow conditions resulting in high recapture probabilities (dashed lines) and high stream flow conditions (dotted lines) resulting in lower recapture probabilities.
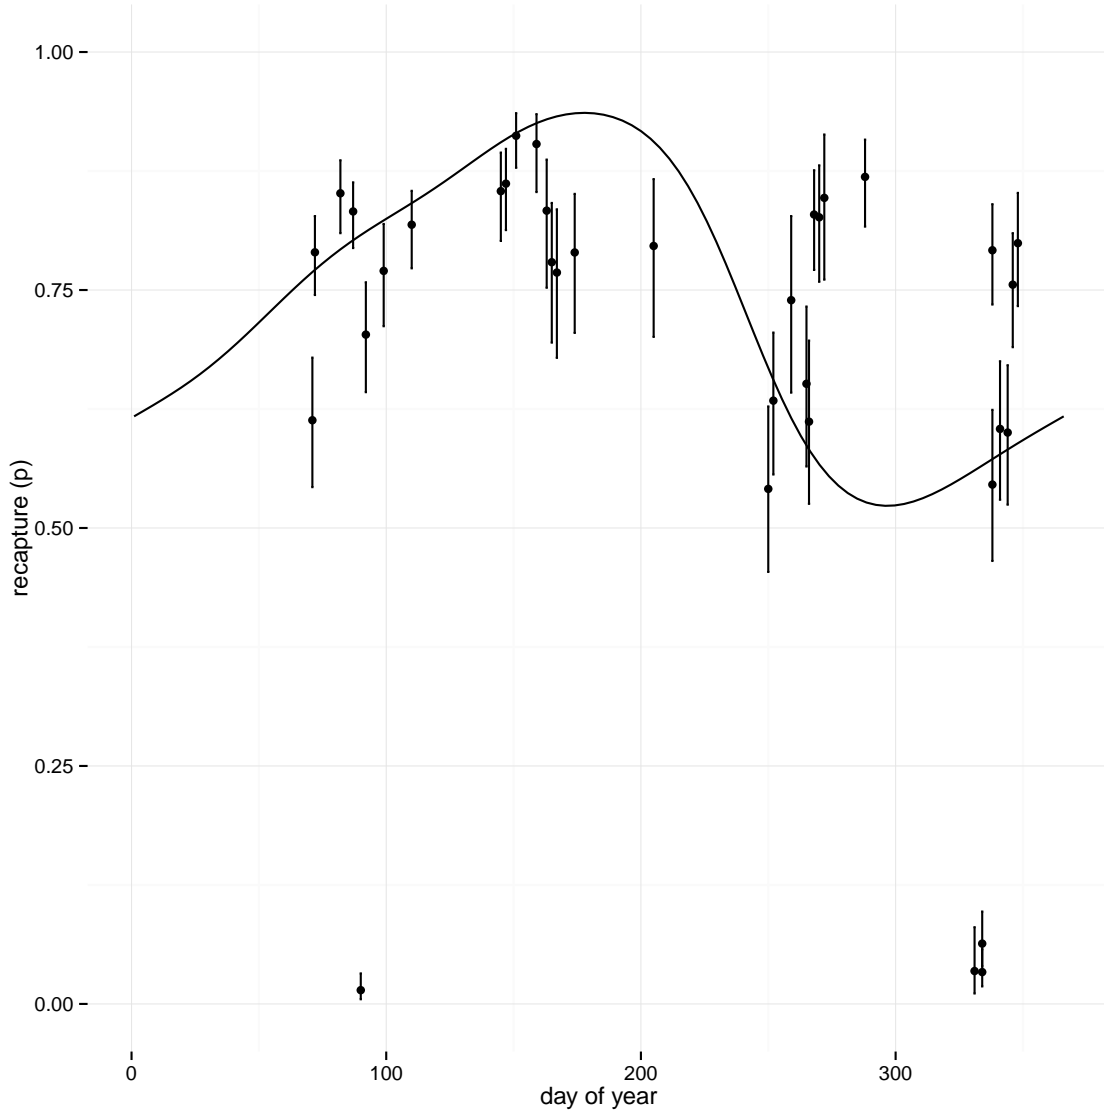
Figure 3: Seasonal and occasion-specific estimates of recapture probability. Results presented here are from the final model. Seasonal-time-specific estimates at average conditions (line) vary smoothly over the year. Sample-specific estimated recapture probabilities (dots are medians with 95% credible intervals) constructed from the seasonal mean plus the per-sample random effect vary widely. Some samples were excluded as the tagged population was not large enough to estimate recapture effectively.
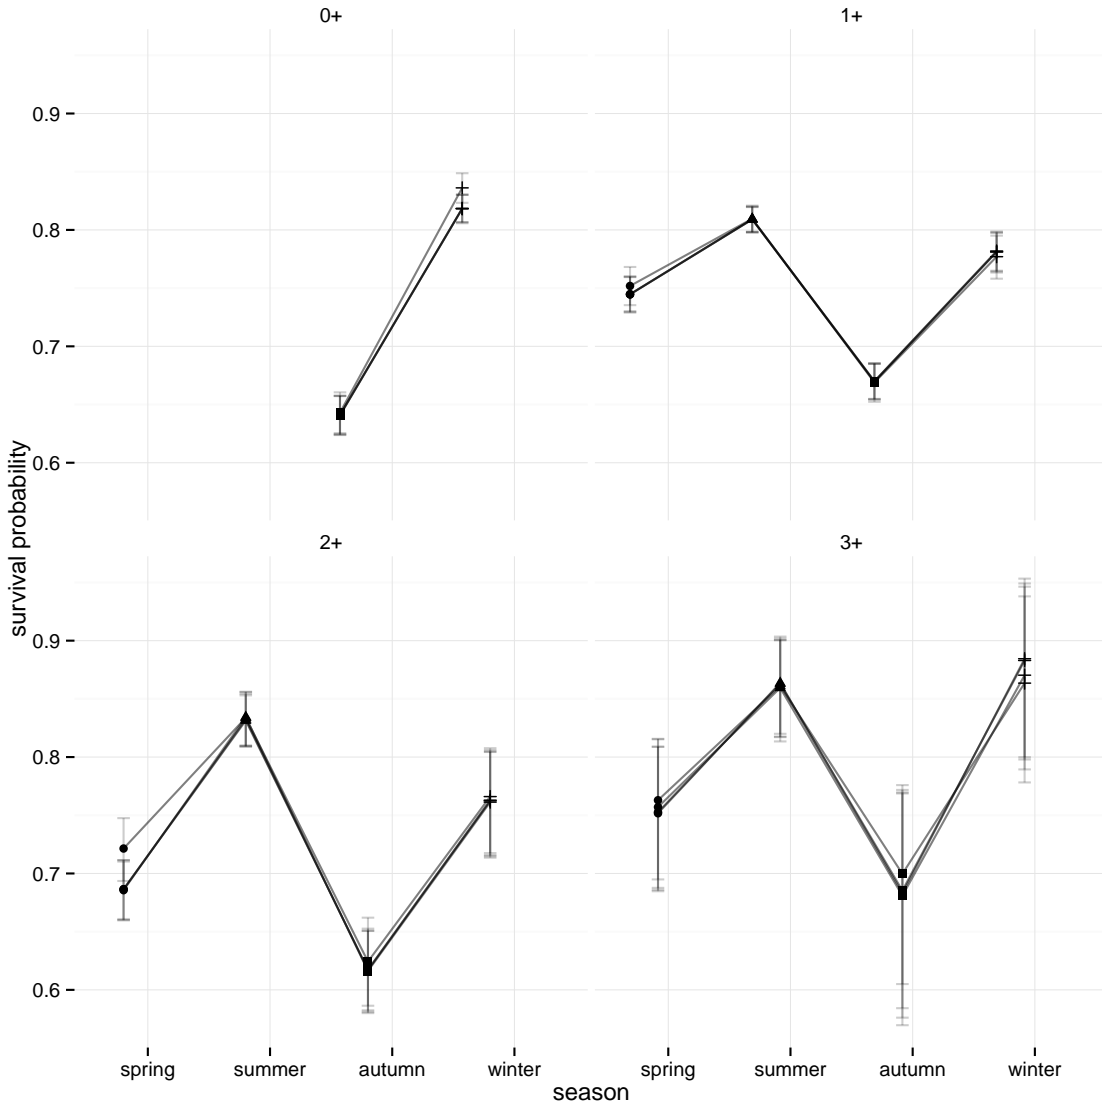
Figure 4: Atlantic salmon survival estimates for each season by age group. The models presented (differentiated by shading) all include seasonal recapture probabilities, as well as environmental effects on recapture. Some include random effects on recapture occasions as well as environmental effects showing that these results are insensitive to the structure of the recapture model.
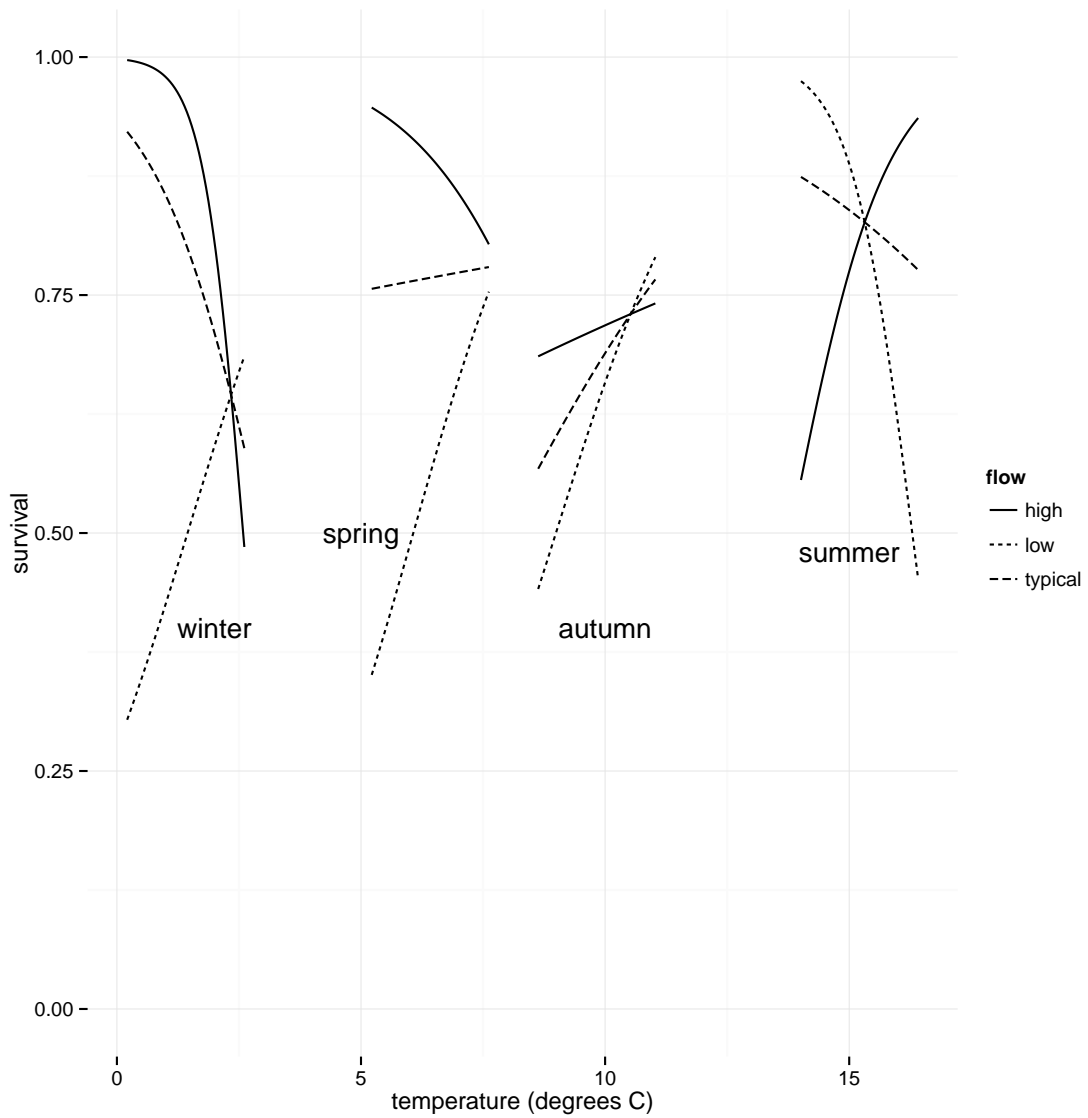
Figure 5: Survival as a function of environmental conditions with flow classified into average, low, and high categories. Low and high categories are defined by offsets of 1.2 standard deviations. Seasons are labeled on the plot and temperature spans from low to high temperature conditions for each season.
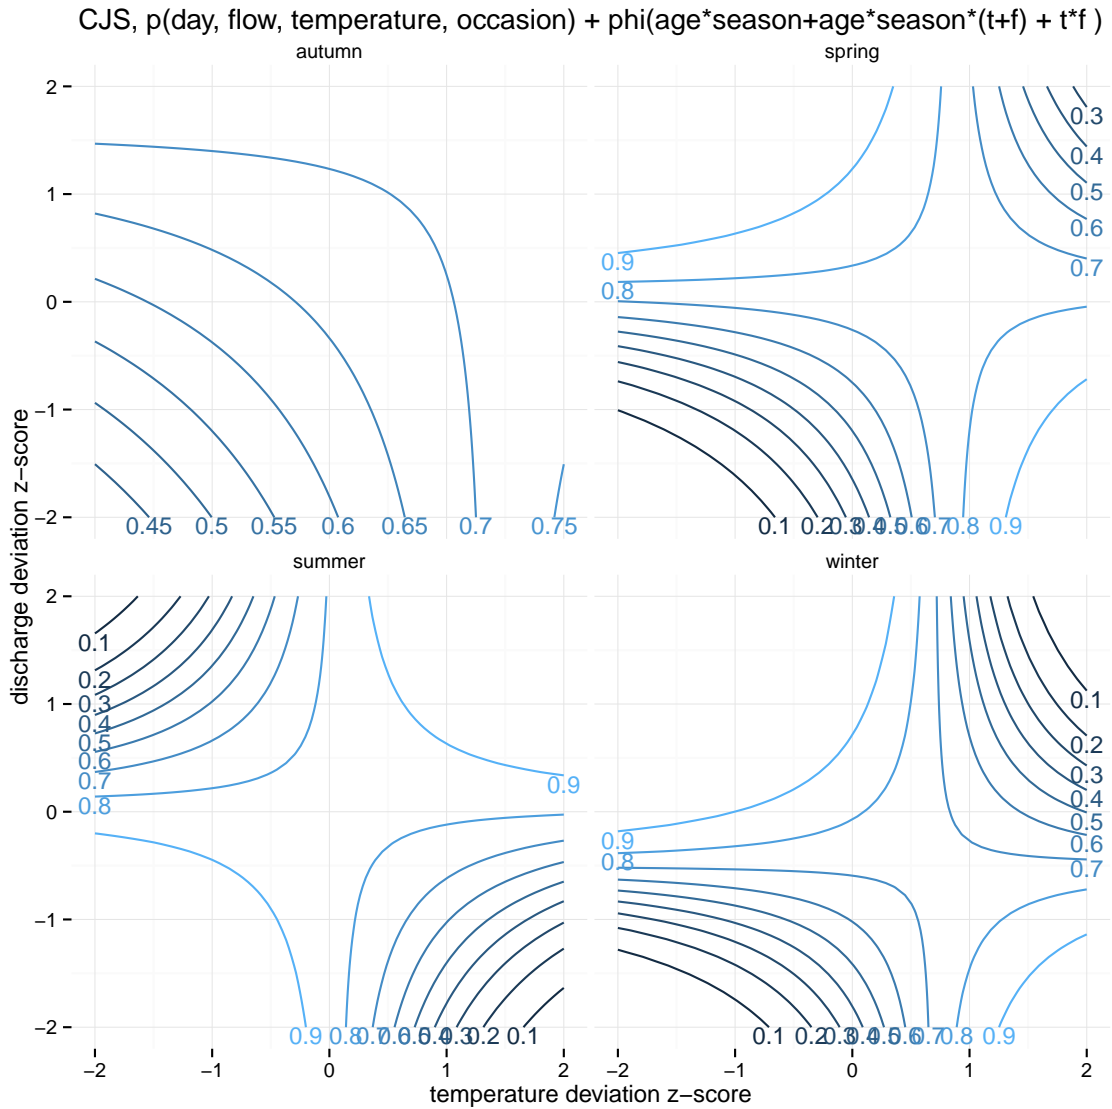
Figure 6: Survival surface as a function of flow and temperature for each season. Result are from the final linear model. Contours are labeled with estimated survival probability at the isoline. Flow and temperature show clear interactions in determining the ultimate survival probability.
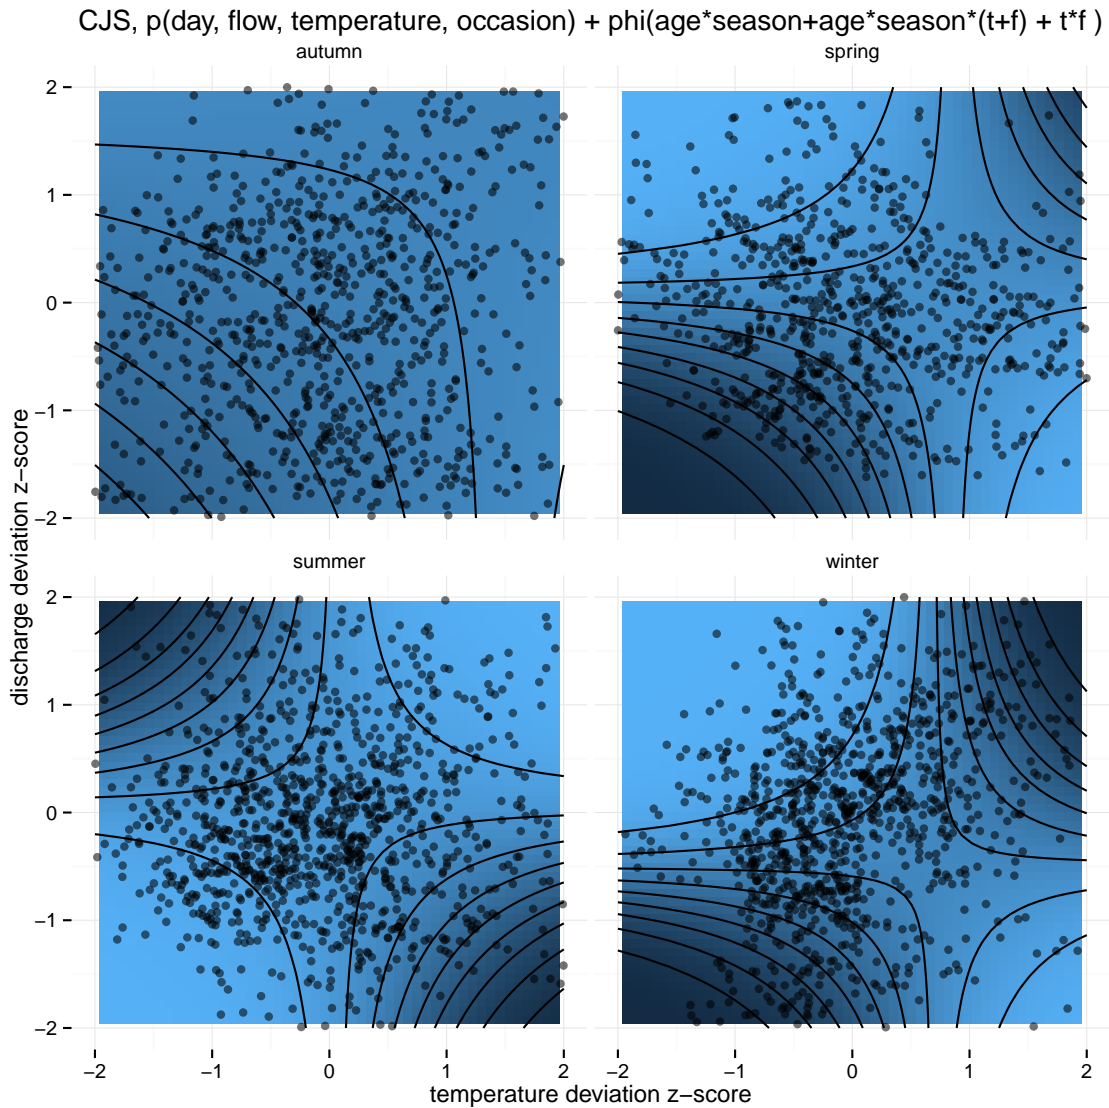
Figure 7: Survival surface as a function of flow and temperature, with data adequacy. The heatmap and a contour plot match the previous contour plot. Areas of high survival are shown in light blue and darker areas indicate low survival. The points represent environmental conditions for all observed season intervals. Effects of flow and temperature vary seasonally and interact to determine their joint effect on vital rates.

Chapter 4

Time-to-event modeling

Discrete data observed at specific time points are common in biological experiments and field studies. These include behavioral choice experiments (Clotfelter *et al.*, 2006), survival experiments (Pollock *et al.*, 1989), visitation experiments (Muenchow, 1986), large scale survival studies (B. H. Letcher, Schueller, *et al.*, 2014), and monitoring using antenna arrays (Zydlewski *et al.*, 2006). A wide range of models is available for modeling such data from logistic regression, survival and reliability analysis, and multiple literatures on multi-state models. Most of these analyses, with the exception of discrete-time multi-state models (Kendall & Nichols, 2002) were developed outside of ecology and experimental biology and only later adapted to common data types in these fields. Survival (reliability) analysis for fully observed data was brought into use in ecology early on by, for example, Muenchow (1986). Continuous-time multi-state models made a more recent entry into the biostatistical literature (Putter *et al.*, 2007). Continuous time models only appear rarely in the ecological literature both because of their relative novelty and the need to adapt them to the realities of field work (Walsh *et al.*, 2015).

Unifying these models under one conceptual framework would be beneficial to the design of experiments and observational field studies. Using a common framework makes it possible to express different models using a common set of parameters. It would encourage direct comparisons of the trade-offs between more or less expensive data types with regards to specific experimental goals.

Sharing a conceptual framework also simplifies analysis, particularly for long-term studies. Many long-term studies experience shifts in methodology due to technological innovations and changing priorities. Without a unified framework it is more difficult to construct analyses of these valuable long-term datasets. Sometimes new technological innovations are disregarded exactly because of ambiguities about how to compare different experimental designs.

I suggest following Muenchow (1986) and the more recent work by Walsh *et al.* (2015) to its roots in the counting process and related theory. This allows the resulting model to handle arbitrary overlapping time scales within a single experiment, and to parameterize multiple otherwise unrelated experiments using a single set of parameters, which is a criti-

cal aspect of combining multiple experimental designs into one analysis, particularly using formal integrated modeling (Schaub, Gimenez, *et al.*, 2007). At the same time these models are able to account for the full range of censoring and other partial observations.

A *counting process*, shown in figure 8, is a way of describing arbitrary events occurring in time. It gives us both a flexible way to describe the timing of these events, and a flexible way of making the event probability calculations required for constructing a likelihood. The process $N(t)$ begins at zero at time t_0 and remains at zero until an event occurs. An event at time t after t_0 causes the process to jump to one. The process then remains at one until another event occurs. On the interval between any two arbitrary time points, defined strictly as $(t_1, t_2]$, the number of events occurring is $N(t_2) - N(t_1)$.
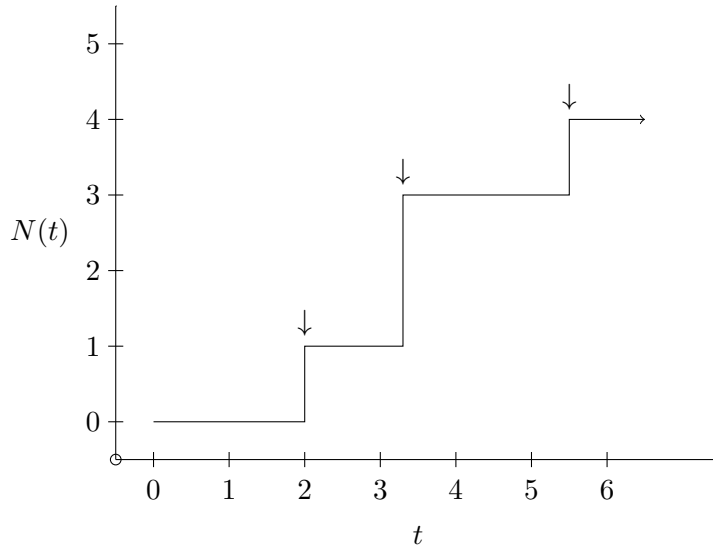
Figure 8: Graphical counting process definition. The counting process is a mathematical formalism useful in both describing the rate at which events arrive, and the probability of a given set of events. Here the counting process $N(t)$ begins at zero at the origin and generates three events indicated by arrows. One event occurs at time 2, two events at time 3.3, and one event at time 5.5 measuring from the origin.

The rate at which events occur in a counting process is defined by the associated *intensity function*, typically written as $\lambda(t)$, which is defined from t_0 onwards and is always either

zero or positive. When $\lambda(t)$ is time-constant, it is often written simply as λ. In all cases it is specified in terms of *events per unit time.*

The focus for inference is typically on the $\lambda(t)$ parameter or function. The expressive power of this model comes from flexible parameterizations for λ and it can produce a number of commonly used models as special cases. When counts are observed, the model can reduce to Poisson regression with offsets. Survival models as well as models from the queue theory literature often use inter-arrival times while observations of point patterns in space are used in kriging. Before returning to the counting process formulation in full, I review closely related models and alternatives, focusing on ecological applications.

4.1   State of the Art

State-change models for discrete data focus on the observed state of identified units. The individual units, indexed by i, are tracked starting at some time $t_{i,0}$ and continuing until either censoring or until they arrive at some immutable state such as death or permanent emigration from the study area. The simplest models track only two states: one original state detected at $t_{i,0}$ and the final immutable state. Figure 9 shows an example with "alive" as the original state and "dead" as the immutable state.

With only two states and a single time period, data such as these fit into the reliability (Hamada *et al.*, 2008) or survival analysis framework (Cox & Oakes, 1984; Freedman, 2008). Specifically, these models consider the probability of event-free survival for the duration d_i, followed by an event (death) at d_i. While many theoretical and applied papers retain the nomenclature of survival analysis regardless of context, "survival" analysis is in practice another word for time-to-event modeling.

Survival models cover a dizzying array of experimental variations including left- and right-censoring, and time-varying covariates. These models cover estimation of absolute survival probabilities as well as proportional hazard estimation that is important in settings where absolute survival might be affected by many complex factors but a small set of isolated factors are of inferential interest. Estimation can be carried out using standard software for survival analysis.

Survival analysis brings to the forefront a set of issues around the choice of reference points for measured times and the bias they can induce. Fieberg & DelGiudice (2009) bring up a variety of ways to anchor and scale time measurements to focus the analysis on
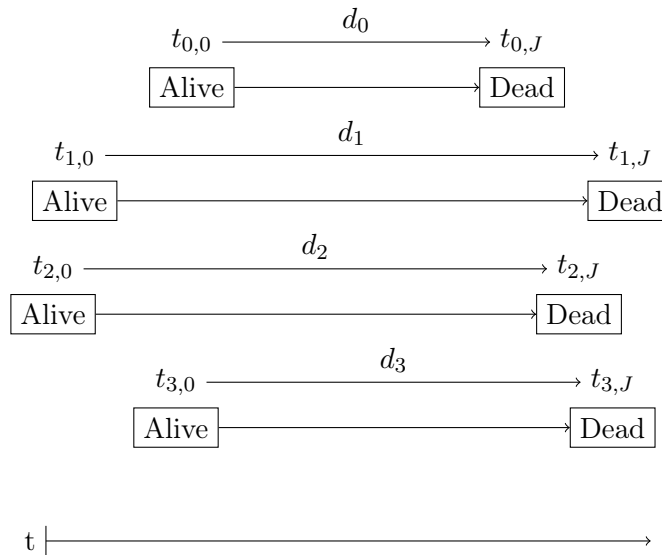
Figure 9: Multi-state model graphical example. Multi-state models focus on the measured time of to transitions between discrete states. The simplest models consider only two states and one time interval, which I show for four units.

seasonal patterns, patterns in time-since treatment, or age-dependent survival. Wolkewitz *et al.* (2010) use the humorous example of actors at-risk for an Oscar, and also at-risk for death after an Oscar to discuss sources of bias in estimates due to mistaken assumptions about what the at-risk pool is.

One lesson of Wolkewitz *et al.* (2010) is that in many simple scenarios and practical experiments, the at-risk pool is more complex than it would otherwise appear and more than one transition needs to be considered. Figure 10 shows the states for a choice experiment where the subject begins the trial at time $t_{i,0}$ and can choose either state A or state B first, potentially followed by the remaining choice. In this case individuals at risk to make choice A include both individuals in the beginning state and individuals who have made choice B. If switching between the two choices was not possible (the choice was an absorbing state) then the risk pool would exclude individuals who made choice B from the at-risk pool for choice A. Complex at-risk pools are better represented by a multi-state model.

Multi-state analyses have gained traction in the medical trial and epidemiological (Ieva
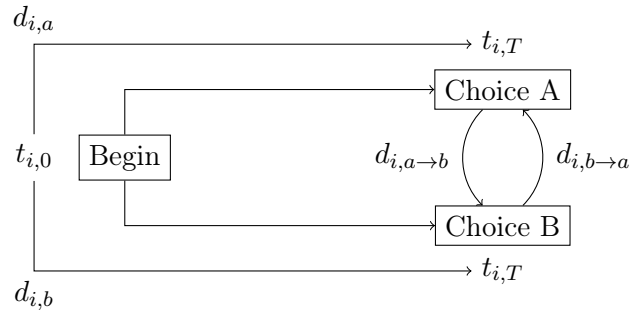
Figure 10: Continuous-time multi-state model graphical example. A more complex multi-state model might consider the time of from the initiation of an experiment to two alternative outcomes, shown here for only a single unit. In some contexts the choice is an absorbing state and only one of the two can be measured on each unit. In other experiments both can be measured and compared.

et al., 2015) literature as an effective way of comparing outcomes for groups of patients who might follow complex trajectories from illness to recovery or some other state. They are supported by two flexible and established packages for the R programming environment—**msm** developed by Jackson (2011), and **mstate** developed by Wreede *et al.* (2011).

While both survival models and continuous-time multi-state models as implemented handle complex state transitions, censoring, as well as individual-specific time-varying covariates, they require fully observed data. Censoring can not account for the case where the state of a unit is unknown at last observation. This is one of the major barriers to applying them to field-based studies where the ability to acquire an observation is uncertain. A second barrier is that event times must typically be observed with high accuracy relative to the time between state changes.

Discrete-time multi-state models are increasingly common in ecology, supported by tools like program **MARK** (White & Burnham, 1999) and standard BUGS language implementations (Lunn *et al.*, 2000; Plummer, 2003). These models provide a rich and generic framework for analysis of discrete state transitions (e.g., stage-based growth, survival, patch-to-patch movement). They have been used to analyze survival, growth, and movement data making them key to estimation of vital rates in field studies.

One assumption of discrete-time multi-state models is that individuals within an experiment must be sampled around the same time (figure 11). This is a loose requirement in the sense that near-equivalent recapture times make the meaning of period-specific state change probabilities less representative of the standard period. It can be managed by discretizing time on a finer scale and estimating state change probabilities over multiple contiguous intervals of time. Both of these work-arounds are only partial as they create ambiguity around the meaning of parameters, multiply the required number of parameters, and therefore scale poorly to complex datasets, which are already burdensome due to number of parameters required.

The assumption of near-equivalent recapture times enforces planning on the logistical side of field studies. In practice it means that the study must be designed to attempt a capture of all individuals in the study population with non-negligible probability. The typical logistical requirement is many people for a short period of time with multiple sets of gear (i.e.,inefficient and difficult to plan). For studies that cover multiple years, these logistical challenges also make it unlikely that seasonal recapture attempts will occur at the same time of the year. Shifting seasonal sampling times make seasonal parameters more difficult to interpret and requires the above-mentioned work-arounds.

Discrete-time models have the advantage of being compatible with many experimental designs for field studies because they typically model the probability of detection in addition to state changes. Units are modeled as observed with a probability, typically labeled p, so when an observation is attempted there is no guarantee that the state will be known. Information about p comes mostly from the proportion of units that are known to be available for observation at time t_j but are unobserved. In turn, if a unit is observed and then never observed again despite repeated attempts, p helps probabilistically bound the time that a unit might stay in its observable state. The last attempted observation time is important because after that time the calculations using p are no longer relevant and the unit is effectively censored as in typical survival models.

The inclusion of detection probability is the main feature that distinguishes the discrete time model structure from the continuous time model structure as available in the literature. Discrete-time likelihood calculations with a detection probability term are straightforward prior and up to the last successful detection time. With event probability θ, and an observation probability p, the likelihood calculations are:
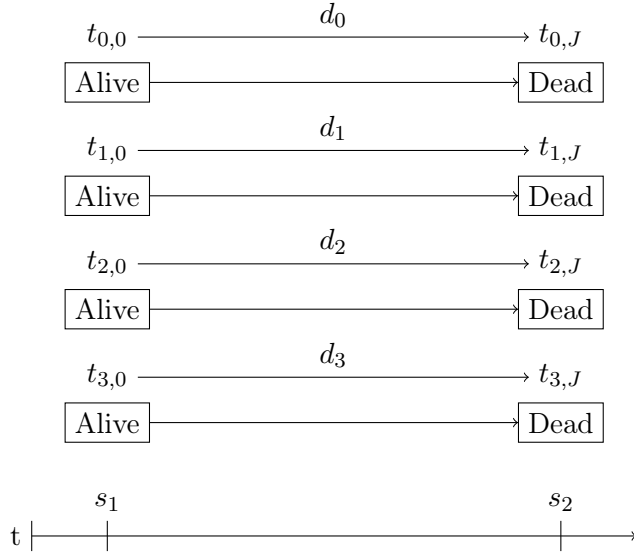
Figure 11: Discrete-time multi-state model graphical example. Discrete-time multi-state models focus on the probability of transition between states between fixed time periods. The simplest models consider only two states and one time interval, which I show for four units.

$$Pr[X = x] = \begin{cases} \theta p, \text{ event occured and state was observed.} \\ \theta(1 - p), \text{ event occured and state was not observed.} \\ (1 - \theta)p, \text{ event did not occur and state was observed.} \\ (1 - \theta)(1 - p), \text{ event did not occur and state was not observed.} \end{cases} \tag{7}$$

As first noted by Darroch (1959), the likelihood calculation after the final detection is more complex but can be carried out using a recursive relationship. The recursive product calculation can be extended up to the last observation attempt and turns into simple censoring with the end of the study. The likelihood calculation for transition to a single absorbing and unobservable state (common in survival models for field data) is:

$$\theta_t + (1 - \theta_t)(1 - p_t)\theta_{t+1} + (1 - \theta_t)(1 - p_t)(1 - \theta_{t+1})(1 - p_{t+1})\theta_{t+2} + \ldots \tag{8}$$

More states and non-observable non-absorbing states make the calculation more complex but conceptually parallel. The summation must extend over the set of possible states that would lead to a non-observation—either removal from the observable pool or a non-detection, at each time step. It is not commonly implemented in the continuous time setting but see Walsh *et al.* (2015) for an implementation with piecewise constant hazard functions and section 4.2.1 for more flexible hazard functions.

The rest of the chapter is organized as follows. First, I state all the necessary details of the counting process required for our calculations and clarify the notation. Second, I introduce discrete-time covariates that can apply to a unit or a group over interval of time but remain constant for the period. Third, I introduce continuous-time covariates that can also apply to a unit or a group over an interval of time but need not maintain a constant effect. This generalization removes the need to rely on piecewise constant hazard functions and as a computational compromise it is implemented using integrable splines. Finally, I discuss the recursive calculation required to implement an observation process model for continuous-time multi-state models but leave the implementation for chapter 3.

4.2 Counting process specification

The data recorded in discrete-state continuous-time experiments are the time of occurrence for multiple types of events. Events might be emitted by a system, such as when worker bees arrive at a hive entrance. Equivalently events might be state changes in one of many units composing the study, such as when an animal gives birth or when a GPS unit stops functioning. The object of interest is the probability of an event—or lack thereof—on a given time interval.

In any case, events of different types for a given unit (or system) can simply be laid out on a time line. We show the point process with its corresponding counting process (figure 12) where the counting process jumps in increments of one at each event. To use this formalism I assume that only one event can ever occur at one point in time. In a continuous time process this restriction is not an important constraint as even co-occurring events can be jittered by a very short time period without significantly affecting calculations.

The Poisson process adds a weak distributional assumption to the counting process. A Poisson process is a counting process $(N(t) : t \geq t_0)$ with the following properties:
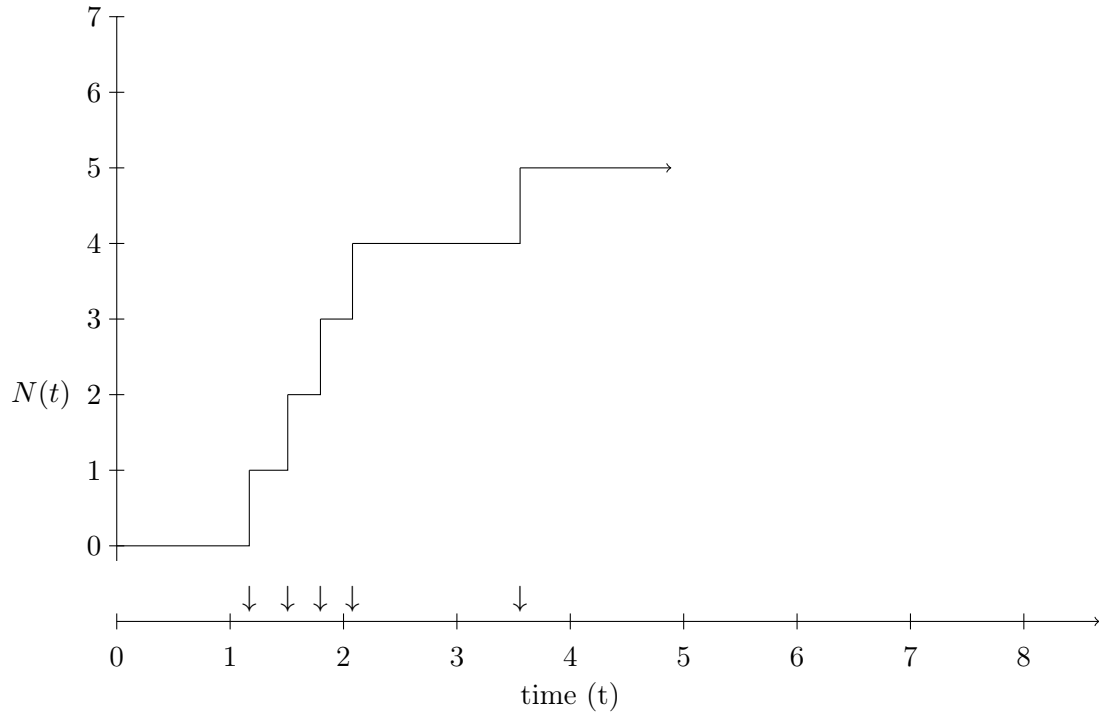
Figure 12: Hypothetical timeline of events in one unit. Both the temporal point process view of discrete-event continuous-time data, and the counting-process view are included. The point process view corresponds to how data is often collected and visualized whereas the counting process view corresponds to a useful formalism.

$$Pr[N(t_0) = 0] = 1 \tag{9}$$

$$N(t_{j+m}) - N(t_j) \sim \text{Poisson}(\Lambda(t_{j+m}) - \Lambda(t_j)) \tag{10}$$

$$\Lambda(t_{j+m}) - \Lambda(t_j) = \int_{t_j}^{t_{j+m}} \lambda(u)du = \Lambda(t_j, t_{j+m}) \tag{11}$$

The properties are 1) the process always begins at zero (at some arbitrary time), 2) between any given timepoints, $t_j < t_{j+m}$, the number of events follows a Poisson distribution; and 3) the expected number of events can be re-written as an integral. When $\lambda(t)$ is

not constant this is referred to as an inhomogeneous Poisson process (IHPP) model or less commonly as a non-homogeneous Poisson process model. The first is a probabilistic statement of the definition of the counting process starting at zero and the second says that *on arbitrary interval*, chosen by the analyst, the number of events follows the Poisson distribution (and its mean-equals-variance assumption). While this assumption is somewhat restrictive, it is no more restrictive than models that use discrete-time event probabilities or piecewise-constant hazard functions as these also imply strict mean-variance relationships. The final property of writing the number of events as an integral provides a flexible way of specifying relationship in continuous time.

4.2.1   Time-varying rate functions

Most continuous-time models in the literature apply piecewise-constant intensity functions to simplify the conversion between the instantaneous event rate, $\lambda(t)$, and the cumulative expected event count, $\Lambda(t_j, t_{j+m})$. I avoid piecewise-constant hazard functions to take advantage of the flexibility of a continuous time model formulation. Now I need to specify $\lambda(t)$ in a broadly applicable way. The typical requirements for $\lambda(t)$ are that the function should be easy to calculate for any given interval, positive everywhere, and that it should be integrable. It is helpful when the function is simple to parameterize.

I suggest that radial splines with any simple-to-integrate basis functions fit these characteristics. Splines can be bounded above zero either by requiring positive coefficients for basis functions or by using a link function. Spline functions are specified as sums of basis functions, which means integration depends only on the basis function and not the complete spline function. Radial basis functions also provide a simple interpretation by relating the knot weight to an approximate local intercept. A radial spline with K knots, centered at knot points μ with scale σ is specified as follows:

$$\lambda(t) = \sum_{k=1}^{K} \beta_k \times f(t, \mu_k, \sigma) \tag{12}$$

Integration relies on the choice of the basis functions, labeled f above. Any integrable basis function would work, but I use the probability density function of the normal distribution (PDF_N) for our basis functions due to their general familiarity and the availability of integrals in most computational packages. The integral of the PDF is the cumulative distri-

bution function used in common theoretical quantile calculations. Due to this property, the relationship between the instantaneous event rate and expected event count is as follows:

$$\Lambda(s,t) = \int_s^t \lambda(x)dx == \int_s^t \sum_{k=1}^K \beta_k \times \text{PDF}_N(x, u_k, \sigma)dx =$$

$$= \sum_{k=1}^K \beta_k \times \int_s^t \text{PDF}_N(x, u_k, \sigma)dx == \sum_{k=1}^K \beta_k \times \text{CDF}_N(x, u_k, \sigma)|_s^t =$$

$$= \sum_{k=1}^K \beta_k \times [\text{CDF}_N(t, u_k, \sigma) - \text{CDF}_N(s, u_k, \sigma)] \quad (13)$$

The flexibility of this function is established either by limiting the number of knots as appropriate to the problem at hand, introducing a penalty term for differences in knot weights or the curvature of the resulting spline, or introducing an equivalent penalty by placing priors on the knot weights. Frequentist approaches are discussed by Wood (2006), and Bayesian formulations are also well established with Eilers & Marx (2010) as the canonical reference. For simplicity I limit the number of knots in our applications as I know the general temporal scale of patterns I am interested in and fine-scale data are often not available in ecological applications.

4.2.2 Event time calculations

The Poisson process model discussed above gives a straightforward way to calculate the likelihood for a flexible intensity function, $\lambda(t)$, and counts of events that arrive in batches where

$$Y \sim \text{IHPP}(\Lambda(t_j, t_{j+k})) \quad (14)$$

Estimating this model requires observations of the batch count, $y_{i,j}$, the start time of counting for the batch t_j, and the end time of counting for the batch, t_{j+k}. This is a common data type and a relatively flexible model.

When arrival times are observed in more detail than batch counts, they can always be summarized to batch counts. However, summarizing them makes it more difficult to

characterize changes in $\lambda(t)$ within batches. Fortunately, it is possible to look at inter-arrival times directly by transforming the count model into an inter-arrival time model that deals more naturally with data at the single event level.

I present a generic IHPP model for time-to-event data below. The random variable for IHPP-distributed event times is T and the realizations of event times are $t_{i,j}$. To complete the model I specify a random variable $Z_{i,j}$ that takes a value of 1 when the event time is observed and a value of 0 when the event is not observed. The probability of observing an event at time $t_{i,j}$ given that it occurs is p and zero otherwise.

$$T_{i.j+k} \sim \text{IHPP}(\Lambda(t_j, t_{j+k})) \tag{15}$$

$$Z_{i,j+k} \sim \text{Bern}(p_{i,j}) \tag{16}$$

The matching indexing between T and Z indicates that a realization of Z is defined only when an event occurs. This is a common all-or-nothing observation process where and observation is acquired with probability p_j if and only if an event occurs.

These random variables have associated realizations $t_{i,1}, \ldots, t_{i,J}$ observed on specified intervals. For a single process in either the univariate or multivariate context we might be interested in likelihoods for any of the following:

- A single event observed at s_2 after s_1.

- A single event observed on the interval $(s_1, s_2]$.

- No event observed on the interval $(s_1, s_2]$.

- No event observed on the interval $(s_1, s_2]$ with imperfect observations.

- A single event not observed on the interval $(s_1, s_2]$, when the probability of detecting events from that interval is p

I discuss each of these in turn below. With these calculations in hand it is possible to compose a model matching any of the event models I review above and integrate inferences from different models using shared parameters.

4.2.3   Univariate likelihood calculations

4.2.3.1   No event observed on the interval $(s_1, s_2]$

In complex studies with time-varying covariates and many discrete periods, the no-event and single-event likelihood calculations are common. The no-event calculation arises when a unit crosses some discrete time boundary—for example the end of a study or an arbitrary boundary between seasons—without experiencing an event. The corresponding likelihood calculation for no events on the interval $[s_1, s_2)$ uses an implied but still unknown density function for event time—$p(t|t > s_1, \lambda(t))$—which can be turned into a direct calculation using the Poisson probability mass function:

$$\ell(\lambda(t)|t > s_1, t > s_2) = \int_{s_2}^{\infty} p(t|t > s_1, \lambda(t)) =$$

$$= Pr[N(s_2) - N(s_1) = 0|\Lambda(s_1, s_2)] = \exp(-\Lambda(s_1, s_2)) \quad (17)$$

The first step restates the desired likelihood as the integral of an unknown density function. The second step transforms this into its corresponding counting process specification, and the final step relies on the Poisson probability mass function producing zero events. I use this expression to derive the rest below.

4.2.3.2   A single event observed on the interval $(s_1, s_2]$

The single-event on an interval calculation arises when an event does occur at an unknown time t on the interval $(s_1, s_2]$. This corresponds to when a state-change occurs during a study and only the state changed is available for observation, not its time. The single-event likelihood calculation can be derived from the no-event starting point. It is worthwhile pointing out that the second step of the probability calculation is much more complex than the rest and I avoid expressing it in detail.

$$\ell(\lambda(t)|t < s_2, t > s_1) = Pr[N(s_2) - N(s_1) > 0|\Lambda(s_1, s_2)] =$$

$$1 - Pr[N(s_2) - N(s_1) = 0|\Lambda(s_1, s_2)] = 1 - \exp(-\Lambda(s_1, s_2)) \quad (18)$$

This likelihood is also the cumulative density function of the first arrival time of the IHPP after s_1, which I can use to define the density function for the first arrival time after s_1.

4.2.3.3 A single event observed at s_2 after s_1

The single-event at time s_2 calculation occurs when the time of the event is observed precisely. The density for a known arrival time can be calculated from the derivative of the CDF from above with respect to the arrival time.

$$\ell(\lambda(t)|t = s_2, t > s_1) = p(t = s_2|\Lambda(s_1, s_2), t > s_1) =$$
$$\frac{d}{dt}\left[1 - \exp(-\Lambda(s_1, s_2))\right] = \exp(\Lambda(s_1, s_2)) \times \left[\frac{d}{dt}\Lambda(s_1, s_2)\right] \quad (19)$$

In the piecewise-constant version of the NHPP model, $\Lambda(s, t)$ is $\lambda \times t$ and the remaining derivative is simply λ following Walsh *et al.* (2015). I maintain a little more generality by assuming that $\Lambda(s, t)$ has a derivative with respect to t, which I write as $\lambda(t)$ following the definition of the properties of the IHPP. In both cases $\lambda(t)$ is a rate with units of events per unit time. If the rate is constant, the number of events on an interval is simply $\lambda \times t$, which produces the correct units:

$$\frac{\# \text{ of events}}{day} \times \# \text{ of days} = \# \text{ of events} \quad (20)$$

The density for inter-arrival times that I need for the single-event likelihood calculation, is then:

$$\ell(\lambda(t)|t = s_2, t > s_1) = \lambda(s_2)\exp(-\Lambda(s_1, s_2)) \quad (21)$$

The single event calculation is sufficient to define the product for a likelihood of a series of successive events, but it becomes inconvenient when sequential arrival times must be integrated out. In that case the direct Poisson model may be more appropriate. An alternative for a set of J events on an interval $(0, t_J]$, appears in Yakovlev *et al.* (2008), which gives an expression for a complete set of events. This expression must typically be calculated using numerical integration but certain forms can be simplified:

$$p(s) = \frac{1}{\Lambda(t_J)} \left[\lambda(s) \exp(-\Lambda(s)) + \int_0^{t_J-s} \lambda(t)\lambda(t+s) \exp(-\Lambda(t,s))dt \right] \tag{22}$$

It considers all possible inter-arrival times for all possible total event counts on the interval. I only include it for completeness and do not apply it further.

4.2.3.4 No event observed on the interval $(s_1, s_2]$ with imperfect observations

The case of partially observed event times arises in the ecological literature more than in areas where continuous-time survival has been applied historically. Up until the final observation, calculations are carried out as above with additional p or $(1-p)$ terms in the likelihood product. After the final observation there are two general cases. The first is only relevant when the event is a change in observable state such that the new state can be observed after the event takes place. Observations are attempted at discrete times with probability of success p and the likelihood of a non-observation is calculated. In the second case, there is a probability p of observing any particular event that occurs at time t.

The first problem is typically solved in discrete-time models following Darroch (1959) using a recursive calculation for observations occurring with probability p at discrete times. The recursive calculation also shows up in the continuous case but only as a final computational step.

I calculate the likelihood as follows. First define the interval when the event occurs as some time in the future after s_1; I write this (s_1, ∞). Second I define our observation data as $x_1, \ldots, x_J$, with associated random variables $X_1, \ldots, X_J$ and the times when the data were collected as $t_1, \ldots, t_J$. The model for the observed data is then Bernoulli with p for the observation probability and an indicator function that is one when the observation times t_j are less than the unknown event time t and zero otherwise.

$$X_j | t_j, t, p_j \sim \text{Bern}(\mathbb{1}_{t_j < t} \times p_j) \tag{23}$$

The likelihood for our non-observations with an unknown event time t after s_1 is then:

$$\ell(\lambda(t), p) = f(x_1, \ldots, x_J | \lambda(t), \vec{p}, s_1) \tag{24}$$

This can be rewritten in terms of an integral for the likelihood with a known event time and all observations in $\vec{x}$ equal to zero:

$$\ell(\lambda(t), \vec{p}) = \int_{s_1}^{\infty} \ell(\lambda(t)|t = s_2, t > s_1)\ell(\vec{p}|\vec{x}, t)dt =$$

$$= \int_{s_1}^{\infty} \lambda(t) \exp(-\Lambda(s_1, t)) \prod_{j=1}^{J}(1 - (\mathbb{1}_{t_j < t} \times p_j)) \times dt \quad (25)$$

All the continuous-time calculations here can be carried out with time boundaries between continuous periods at observation times (or other arbitrary boundaries) since no event is ever directly observed. Ultimately the integral here must be computed in a piece-wise fashion due to the indicator functions and an example appears in equation 25.

4.2.3.5 A single event observed at s_2 when the probability of observing a given event is p

A second case of partially observed events occurs when observations are continuous but can only be made successfully with probability p when an event occurs. If no auxiliary data are available, the observation parameter is fully confounded with the observation probability leading to $p\lambda(t)$ as the inferred rate function. If auxiliary data are available, model-based inference of p can be used to reduce confounding. The model-specific discussion is beyond the scope of this section.

4.2.4 Multivariate likelihood calculations

Multivariate time-to-event calculations follow the same patterns as the univariate calculations but also consider the issue of competing risks. With K competing processes, an event from processes 1 at time t means that no events occurred from competing processes. Exact calculations can still be relatively straightforward when t is observed but become difficult when t is unobserved. In all cases I consider K competing independent processes. For each independent process, the distribution of inter-arrival times follows the inhomogeneous Poisson process distribution as above with the random variables T_1 through T_K and the dummy variables t_1 through t_K with the first-event time density functions:

$$p_1(t) = \lambda_1(t_1) \exp(-\Lambda_1(0, t_1)) \tag{26}$$

$$\ldots \tag{27}$$

$$p_K(t) = \lambda_K(t_K) \exp(-\Lambda_K(0, t_K)) \tag{28}$$

Due to independence, the joint distribution is simply the product: $p(t_1, t_2, \ldots, t_K) = p_1(t_1)p_2(t_2)\ldots p_K(t_K)$. The calculations below are stated as integrals over this joint distribution. Where I can, I maintain generality by using $\lambda(t)$ and $\Lambda(s, t)$ but in cases where only the homogeneous case can be simplified, I assume homogeneity and use λ and $\lambda \times (t - s)$.

4.2.4.1 No event observed on the interval $(s_1, s_2]$

When no event occurs on $(s_1, s_2]$, it is known that none of the processes produced an event, or equivalently the probability that the event occurred after time t_2. The corresponding calculation, starting from the joint distribution and following the definition of a no-event likelihood derived above is:

$$\ell(\lambda(t_1), \ldots, \lambda(t_K)|t_1 > s_2, \ldots, t_K > s_2, t_1 > s_1, \ldots, t_K > s_1) =$$

$$Pr[t_1 > s_2, \ldots, t_K > s_2 | t_1 > s_1, \ldots, t_K > s_1] =$$

$$\int_{s_2}^{\infty} \cdots \int_{s_2}^{\infty} \prod_{k=1}^{K} p_k(t_k|s > 1, \lambda(t_k)) = \prod_{k=1}^{K} \exp(-\Lambda(s_1, s_2)) \tag{29}$$

When no events occur there are no interactions among the K processes and the joint likelihood is simply the product of the univariate likelihoods.

4.2.4.2 A single event observed on the interval $(s_1, s_2]$

When an event occurs somewhere on $(s_1, s_2]$, the calculation is complicated by competing risks. Assuming the event from process j occurs at time t_j, we learn about the outcomes of a variety of processes. Namely process j produced a single event and failed to produce any further events, whereas the remaining processes produced no events. I start from the joint density for single events, $p(t_1, \ldots, t_K|t_1 > s_1, \ldots, t_K > s_1)$, specified as the product of individual process densities. The interval is divided into two time periods for process j so

63

it enters the product twice, once anchored at s_1 and once anchored at t_j. I integrate over $(s_1, s_2]$ for the no-event processes, and I integrate from t_1 to s_2 for the single event process. Unfortunately I do not know where on $(s_1, s_2]$ the point t_1 is found, so I integrate t_1 over the interval as a final step.

$$\ell(\lambda_1(t_1), \ldots, \lambda_K(t_K) | t_1 > s_2, \ldots, t_j < s_2, \ldots, t_K > s_2, t_1 > s_1, \ldots, t_K > s_1) =$$

$$\int_{s_1}^{s_2} \left[p_j(t = t_j | s_1) dt \int_{t_j}^{s_2} p_j(t | t_j) dt \int_{s_1}^{s_2} \cdots \int_{s_1}^{s_2} \prod_{k=1, k \neq j}^{K} p_k(t) dt \right] dt_j =$$

$$\int_{s_1}^{s_2} \left[\lambda(t_j) \exp(-\Lambda_j(s_1, t_j)) \exp(-\Lambda_j(t_j, s_2)) \prod_{k=1, k \neq j}^{K} \exp(-\Lambda_k(s_1, s_2)) \right] dt_j =$$

$$\prod_{k=1, k \neq j}^{K} [\exp(-\Lambda_k(s_1, s_2))] \times \exp(-\Lambda_j(s_1, s_2)) \times \int_{s_1}^{s_2} \lambda_j(t_j) dt_j =$$

$$\prod_{k=1, k \neq j}^{K} [\exp(-\Lambda_k(s_1, s_2))] \exp(-\Lambda_j(s_1, s_2)) \Lambda_j(s_1, s_2) \quad (30)$$

In certain situations, particularly when events are state changes, the processes active before t_j are not necessarily the same as after t_j that leads to a similar set up but does not simplify as far. I start again from the joint density but this time I divide it into events before t_j when processes from set A are active and events after t_j when processes from set B are active. For a known t_j the likelihoods for the two sets are:

$$\ell(\lambda_1(t_1), \ldots, \lambda_K(t_K) | t_{i \in A} > t_j, t_j = t, t_{i \in B} > s_2, t_{i \in A} > s_1, t_{i \in B} > t_j) =$$

$$p_j(t_j = t | s_1) \int_{s_1}^{t_j} \prod_{k \neq j, k \in A} p_k(t) dt \times \int_{t_j}^{s_2} \prod_{k \in B} p_k(t) dt \quad (31)$$

Since t_j is unknown and present somewhere between s_1 and s_2, the final expression integrates over the possible values of t_j.

64

$$\ell(\lambda_1(t_1), \ldots, \lambda_K(t_K)|t_1 > s_2, \ldots, t_j < s_2, \ldots, t_K > s_2, t_1 > s_1, t_K > s_1) =$$

$$Pr[t_j > s_2|t_j > s_1, \ldots, t_j < s_2|t_j > s_1, \ldots, t_K > s_2|t_K > s_1] =$$

$$\int_{s_1}^{s_2} p_j(t = t_j|s_1) \int_{s_1}^{t_j} \prod_{k \neq j, k \in A} p_k(t)dt \times \int_{t_j}^{s_2} \prod_{k \in B} p_k(t)dt \quad (32)$$

Simplification would typically be model-specific or require a homogeneous Poisson process assumption.

4.2.4.3 A single event observed at s_2 after s_1

A single event likelihood at a known time can be calculate as part 1 from above with a known t_j:

$$\ell(\lambda_1(t_1), \ldots, \lambda_K(t_K)|t_1 > s_2, \ldots, t_j = s_2, \ldots, t_K > s_2, t_1 > s_1, t_K > s_1) =$$

$$Pr[t_1 > s_2|t_1 > s_1, \ldots, t_j = s_2|t_j > s_1, \ldots, t_K > s_2|t_K > s_1] =$$

$$p_j(t = t_j|s_1) \times \int_{s_2}^{\infty} \prod_{k \neq j}^{K} p_k(t)dt = \lambda_j(t_j)\exp(-\Lambda_j(s_1, t_j)) \times \prod_{k \neq j}^{K} \exp(-\Lambda_k(s_1, s_2)) \quad (33)$$

4.2.4.4 No event observed on the interval $(s_1, s_2]$ with imperfect observations

Sometimes the events are not observed, but the state can be observed, before the event only, with a fixed probability. If only a single event is possible (e.g., transition to an absorbing state) the likelihood is much like the univariate form discussed above.

$$\ell(\lambda_1, \ldots, \lambda_K, \vec{p}|x_1 = 0, \ldots, x_K = 0, t_1 > s_1, t_K > s_1) =$$

$$\sum_{j=1}^{K} \int_{s_1}^{\infty} \lambda(t)\exp(-\Lambda_j(s_1, t)) \prod_{j \in (s_1, t]} (1 - p_j) \prod_{k=1, k \neq j}^{K} \exp(-\Lambda_k(s_1, t))dt \quad (34)$$

In practice the integral must be calculated piecewise using the components defined above. For an example see equation 25. When the states that lead to the end of detectability are not absorbing states the sum must cover all possible ways that non-detectability can arise and I do not cover that case here.

4.3 Behavioral choice example

I illustrate a simple three-state model with a behavioral experiment on bluenose shiner carried out in order to understand their nest choice. The goal of this example is to illustrate the continuous-time likelihood calculations on a simple multi-variate model with competing states. I simplify this example by considering only a homogeneous Poisson process with independent rates for different events. The complexity of the calculation comes from calculating competing forward and reverse transitions on the intervals.

A simpler discrete-time approach might model the probability of each transition over a discrete period of time but that construction makes it awkward to consider competing transitions, particularly reverse transitions. Intuitively, a high rate of reverse transitions makes it more likely that when a transition from A to B is observed between s_1 and s_2, the transition occurred closer to s_2 than it might otherwise.

4.3.1 Background and methods

The bluenose shiner is a freshwater fish with an interesting egg-laying behavior centered on the longear sunfish. Longear sunfish are significantly larger than the shiners and male sunfish build and defend a nest. Shiners search for sunfish nests and lay their eggs alongside sunfish eggs with no evidence of a parasitic relationship. The relationship appears to be an example of a mutualistic nest association. Shiners benefit from nest defense by the larger sunfish and sunfish potentially benefit by having the effect of any successful predation diluted by the shiners' eggs. There is also some evidence (*Adam Fuller, personal communication*) that sunfish with longer opercular flaps ("ears") perform better at nest defense.

Fuller & Earley (2015) carried out an experiment to test the hypothesis that shiners in search of a nest location would prefer the nest of a sunfish with a longer opercular flap. For each experimental trial, a 120 cm long tank with five compartments was prepared by placing sunfish in either one or both end compartments (30 cm long by 30 cm wide by 30 cm high). The end compartments were separated from the rest of the tank by a mesh which was too small for the sunfish to pass through but large enough for shiners to move through freely. The sunfish were experimentally manipulated to have one of three opercular flap lengths—short, medium, or long. This ultimately resulted in five experimental treatments defined by the presence of one or two sunfish and their respective opercular flap length(s)—

long vs. medium, long vs. short, long vs. empty, medium vs. short, and short vs. empty. The three middle compartments were left empty. To prepare a trial, typical reproductive groups of two male and one female shiners were placed in the central compartment and kept there with removable partitions. To start a trial the partitions were removed and the movement of the shiners were videotaped. Rather than using shiner egg-laying behavior as the endpoint, videotaping continued for 30 minutes from the start of the trial. Based on the video, the locations but not identities of the three shiners were recorded every 20 seconds. Details about randomization, acclimatization, and other experimental details can be found in Fuller & Earley (2015).

The goal of the experiment was to use location, specifically movement towards sunfish with longer ear flaps, as a proxy for more complex nest preferences. There are a number of difficulties with this experimental design. First, it is unclear whether preference for one side of the tank correlates with egg-laying behavior. Second, the object of interest is the individual behavior but individuals are not identified in the data, rather three sunfish position are identified purely by ordering: left, center, and right. Third, the "choices" available to the shiners were not final. Shiners might quickly move to one side and, having made an internal choice of nesting location, might move away.

The difficulty with interpretation of a non-final choice of location as a proxy for a more complex behavior is beyond the scope of this analysis. I address the difficulty of using locations without identities by evaluating the correlation of locations among the three shiners. High correlation would suggest that a single summary of group location might contain adequate information for an analysis whereas low correlation would mean more complex group dynamics were present. I address the non-final nature of the available choices by focusing the model on the time to transitions between discrete segments of the tank.

In the following analysis, tracks of summary statistics are presented rather than tracks of individuals. The decision to explore the behavior of well-defined summary statistics comes from the fact that individual identities are not known. I use the order statistic of shiner positions (left-most shiner, center shiner, right-most shiner) as well as the group median (effectively the position of the center shiner).

Intermediate positions in the tank might be meaningful, but we can simplify the analysis by focusing on three states—which we label as unsure (U), left (L), or right (R). Shiners were assigned to the left (L) state when they were in the leftmost 30 cm of the tank, they were assigned to the right (R) state when they were in the rightmost 30 cm of the tank, and they

were assigned to the unsure (U) state otherwise. I consider the time to transition between any of these states to summarize a given experiment and we compare estimates across groups to derive inference about preferences. The resulting state and transition duration diagram is shown below in figure 13.
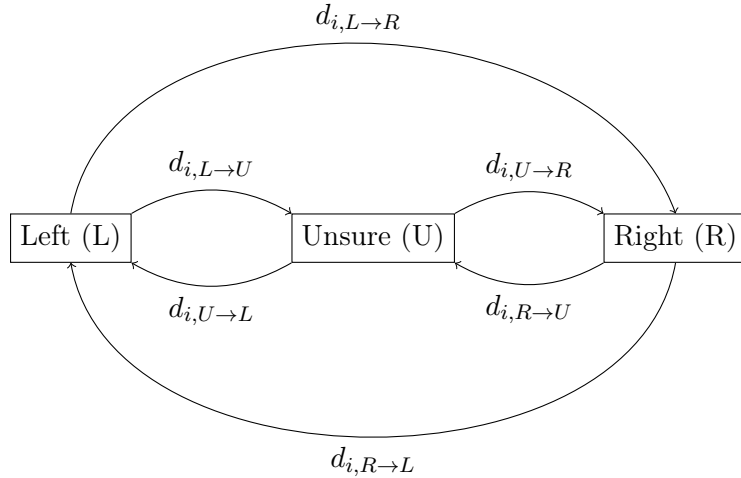
Figure 13: Three states of the shiner experiment. The states are: left (L, leftmost 30 cm of the tank), right (R, rightmost 30 cm of the tank), or unsure (U, elsewhere), match the concept of "choice by location". The observational data can be converted into a time-to-event format (where an event is a state transition) and the probability of a transition to either choice can be modeled and compared across groups to produce inferences about the biological process.

4.3.2 Exploratory analysis

One important choice in analyzing this data set is whether to analyze all three location or whether it is possible to characterize group position overall and analyze the group-level decision making. Analyzing the experiments at the group level is more compatible with the conceptual framework of the experiment and simpler so it is the preferred choice if feasible.

To assess whether a group-level observable can be generated, we plotted the distance of each individual from the bottom left corner of the tank over time. In general shiner loca-

Table 4: Correlations among shiner locations by treatment group.

treatment	LHS corr.	RHS corr.
long vs. medium	0.87	0.95
long vs. short	0.80	0.79
long vs. empty	0.86	0.81
medium vs. short	0.77	0.85
short vs. empty	0.84	0.83

tions were highly correlated to the point where analyzing separate tracks would yield little additional information. Some experiments show high correlations among order statistics (figure 14, top two panels). A smaller number of shiner groups also displayed fusion-fission dynamics (figure 14, bottom left panel) which occur in many group-living species (Conradt & Roper, 2000) but are present here primarily as a nuisance feature of the data. Fusion-fission dynamics make it more difficult to characterize the group choice, particularly with a very small group. Finally, a few groups exhibited more complex group behavior such as frequent splitting or an individual displaying limited movement while being visited by the others (figure 14, bottom right panel). I show a limited set of plots here but I generated and inspected plots for all experiments to assess the viability of a group location definition.

Overall the shiner movements are highly correlated within the group. I calculated Spearman correlation in the distance to the lower left corner by video. The plot by experimental group of these correlations is shown in figure 15 and generally shows high correlation with no apparent difference based on the side of the tank with a larger sunfish.

I also calculated quantiles of the group-wise correlations which are shown in table 4. While there is a substantial amount of variation in correlation among videos, the different treatments do not appear to have a strong effect on in-group correlation. This makes us confident that a choice of group summary would not bias a comparison among treatments.

Based on these summaries and visualizations we choose the central shiner as our summary of the group location. In cases when the group position is highly correlated this choice does not matter much. In cases where the group splits, at least two fish tend to stay together and the middle fish will track the pair. In no experiment does the group split entirely, though some groups do show significant movement over short periods of time and

our summary may not capture this movement effectively.

4.3.3   Model specification

Since the experiment is observed regularly, we have regular observations of the time t, and the state $X(t)$ which indicates where the focal fish is located. Since the process for arriving at a state is interesting, rather than the state itself, we model the time between state-to-state transitions.

A variety of models might be appropriate for this dataset. We focus somewhat artificially on a counting process model with a constant per-transition intensity function with parameter λ. The result of picking this intensity function is an implied exponential distribution for transition times, which simplifies calculations for this multi-variate problem. We estimate λ parameters for each non-self transition in figure 13 and each experimental treatment.

As discussed above we choose the middle shiner to summarize group position. To generate time-to-event data, we classify the position of this shiner as either left (L), unsure (U), or right (R) states. We make an approximate likelihood calculation. In continuous time, for an observed transition U $\to$ R, the exact likelihood is a sum of the likelihood for one, two, or more transitions. We truncate this calculation at the single transition likelihood plus the two-transition likelihood due to the relatively frequent observations and limited likelihood contribution of the many-transition likelihoods.

The transition to self is simplest. A transition to self is modeled as a non-event with regards to the other available other transitions. For example, the likelihood for the "transition" $U \to U$ on the interval $(s_1, s_2]$ is:

$$Pr[t_{\mathrm{U}\to\mathrm{R}} > s_2, t_{\mathrm{U}\to\mathrm{L}} > s_2 | t_{\mathrm{U}\to\mathrm{R}} > s_1, t_{\mathrm{U}\to\mathrm{L}} > s_1] = \exp\left[-(\lambda_{\mathrm{U}\to\mathrm{R}} + \lambda_{\mathrm{U}\to\mathrm{L}}) \times (s_2 - s_1)\right]$$

(35)

A transition from one state to another is more complex as the first term includes a competing risk of the alternative transition in a three-state system but for constant λ parameters and one event on the interval $(s_1, s_2]$, it can be written as the single-event on an interval probability from above, multiplied by the probability of no competing event and no reverse events. We start from the desired probability statement for all t's and restate it in terms of the joint density for independent $p(t_{\mathrm{U}\to\mathrm{R}}, t_{\mathrm{U}\to\mathrm{L}}, t_{\mathrm{R}\to\mathrm{U}}, t_{\mathrm{R}\to\mathrm{L}})$:

$$Pr[t_{\text{U} \to \text{R}} < s_2, t_{\text{U} \to \text{L}} > t_{\text{U} \to \text{R}}, t_{\text{R} \to \text{U}} > (s_2 - t_{\text{U} \to \text{R}}), t_{\text{R} \to \text{L}} > (s_2 - t_{\text{U} \to \text{R}})|$$

$$t_{\text{U} \to \text{R}} > s_1, t_{\text{U} \to \text{L}} > s_1, \lambda_{\text{U} \to \text{R}}, \lambda_{\text{U} \to \text{L}}, \lambda_{\text{R} \to \text{U}}, \lambda_{\text{R} \to \text{L}}] =$$

$$\int_{s_1}^{s_2} \int_{t_{\text{U} \to \text{R}}}^{\infty} \int_{s_2 - t_{\text{U} \to \text{R}}}^{\infty} \int_{s_2 - t_{\text{U} \to \text{R}}}^{\infty}$$

$$p(t_{\text{U} \to \text{R}}|t_{\text{U} \to \text{R}} > s_1, \lambda_{\text{U} \to \text{R}}) p(t_{\text{U} \to \text{L}}|t_{\text{U} \to \text{L}} > s_1, \lambda_{\text{U} \to \text{L}}) p(t_{\text{R} \to \text{U}}, t_{\text{R} \to \text{L}}|\lambda_{rtu}, \lambda_{rtl}) dt_{\text{R} \to \text{U}} dt_{\text{R} \to \text{L}} dt_{\text{U} \to \text{L}} dt_{\text{U} \to \text{R}} =$$

$$\frac{\lambda_{\text{U} \to \text{R}} \exp(-(\lambda_{\text{R} \to \text{U}} + \lambda_{\text{R} \to \text{L}})s_2) \exp((\lambda_{\text{U} \to \text{R}} + \lambda_{\text{U} \to \text{L}})s_1)}{-1(\lambda_{\text{U} \to \text{R}} + \lambda_{\text{U} \to \text{L}} - \lambda_{\text{R} \to \text{U}} - \lambda_{\text{R} \to \text{L}})} \times$$

$$\left[\exp(-(\lambda_{\text{U} \to \text{R}} + \lambda_{\text{U} \to \text{L}} - \lambda_{\text{R} \to \text{U}} - \lambda_{\text{R} \to \text{L}})s_2) - \exp(-(\lambda_{\text{U} \to \text{R}} + \lambda_{\text{U} \to \text{L}} - \lambda_{\text{R} \to \text{U}} - \lambda_{\text{R} \to \text{L}})s_1) \right]$$

$$(36)$$

In general this can be written using the rate function for the event, λ_E, the rate function for forward events, λ_F, and the rate function for the reverse events, λ_R where both λ_F and λ_R sum their component rate functions.

$$\frac{\lambda_E \exp(-(\lambda_R)s_2) \exp((\lambda_F)s_1)}{-1(\lambda_F - \lambda_R)} \times \left[\exp(-(\lambda_F - \lambda_R)s_2) - \exp(-(\lambda_F - \lambda_R)s_1) \right] \quad (37)$$

In the corner case where $\lambda_{\text{U} \to \text{R}} + \lambda_{\text{U} \to \text{L}}$ equals $\lambda_{\text{R} \to \text{U}} + \lambda_{\text{R} \to \text{L}}$ the above expression is undefined and taking a different path to the integral results a simpler expression:

$$Pr[t_{\text{U} \to \text{R}} < s_2, t_{\text{U} \to \text{L}} > s_2, t_{\text{R} \to \text{U}} > (s_2 - t_{\text{U} \to \text{R}}), t_{\text{R} \to \text{L}} > (s_2 - t_{\text{U} \to \text{R}})|$$

$$t_{\text{U} \to \text{R}} > s_1, t_{\text{U} \to \text{L}} > s_1, \lambda_{\text{U} \to \text{R}}, \lambda_{\text{U} \to \text{L}}, \lambda_{\text{R} \to \text{U}}, \lambda_{\text{R} \to \text{L}}] =$$

$$\lambda_{\text{U} \to \text{R}} \exp(-(\lambda_{\text{R} \to \text{U}} + \lambda_{\text{R} \to \text{L}})s_2) \exp((\lambda_{\text{U} \to \text{R}} + \lambda_{\text{U} \to \text{L}})s_1) \times (s_2 - s_1) \quad (38)$$

In general it is possible to have multiple events within a single interval. The general calculation follows the pattern above but may require numerical integration and is beyond the scope of this example. Instead of attempting the entire calculation, we instead assume that the system is observed frequently enough that multiple transitions make a negligible contribution to the likelihood.

The full table of possible transitions and their event likelihood contributions are shown in table 5 below.
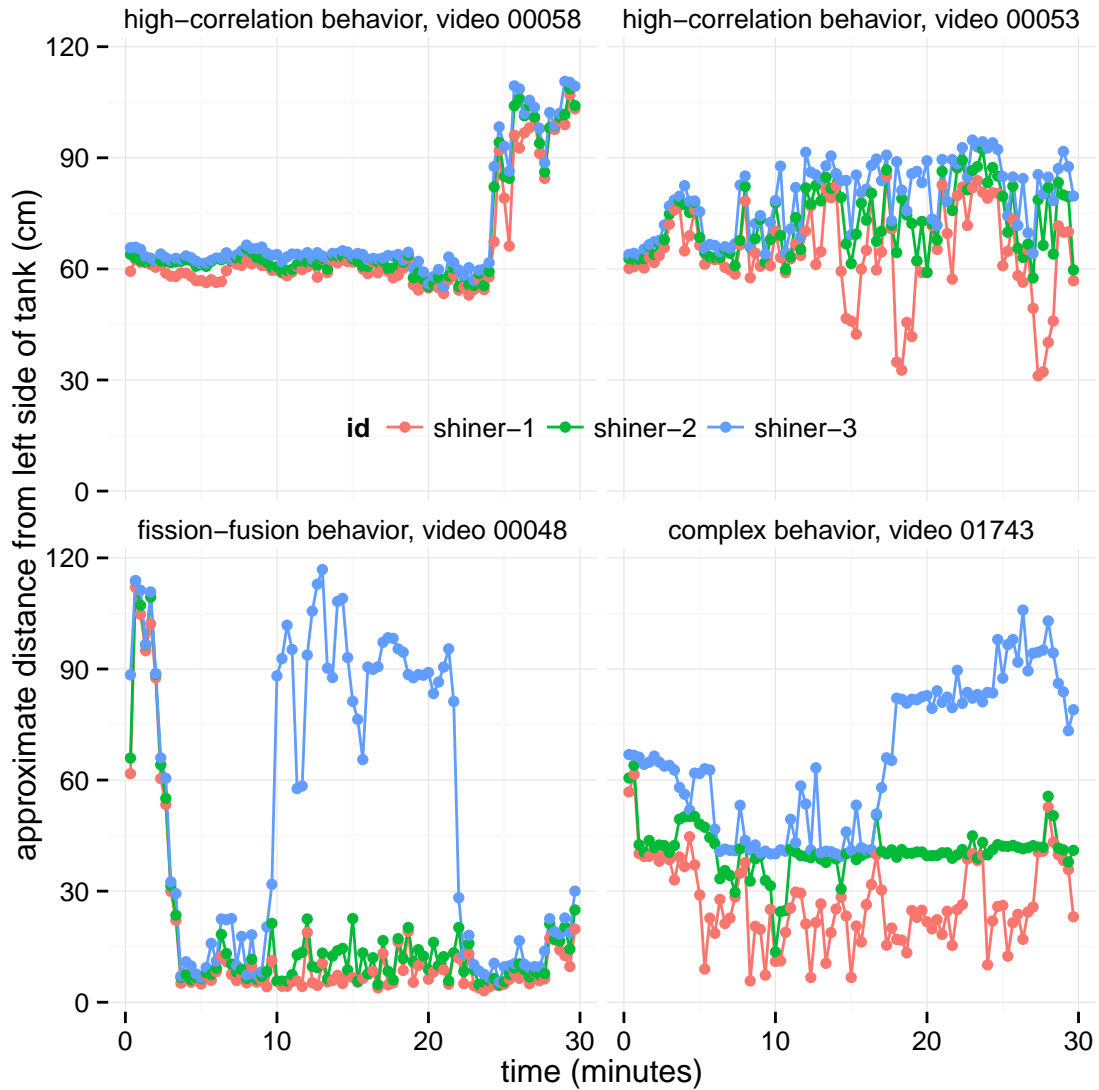
Figure 14: Tracks of shiner position order statistics as distance from the left hand side of the tank over time. Shiner group behavior was typically straightforward with all three individuals forming a tight group (top two panels). Some groups displayed limited fission-fusion dynamics (bottom left panel) and some groups displayed much more complex behavior (bottom right panel) which would be difficult to interpret as any sort of preference for one side or the other.
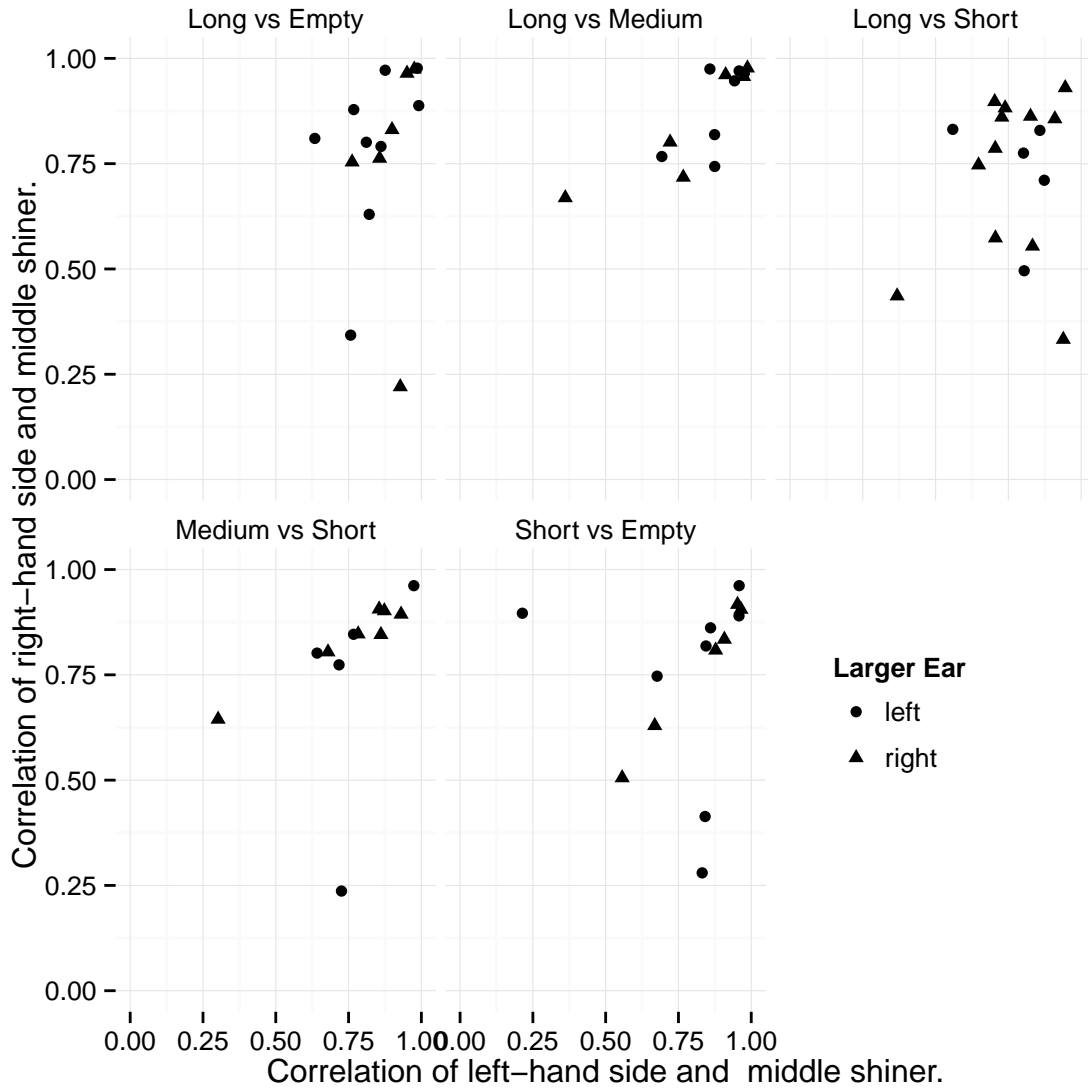
Figure 15: d

istance from the bottom left corner of the tank by treatment group.]Experiment-wise correlation of location among shiners within a group. Position was measured as distance from the bottom left corner of the tank. The x-axis is correlation of the left hand side and middle fish, and the y-axis is the correlation of the right hand side and middle fish. Correlations are generally high with $> 54\%$ of experiments having a minimum correlation above .75

Table 5: Per-event likelihood for the shiner movement data set. The states are: $S(t_i, t_{i+1}) : S \in L, U, R.$ and a transition is defined by a pair of states.

Transition No.	Transition	$\ell(T\|\lambda_i)$
1	$U \to U$	$\exp\left[-(\lambda_{U \to R} + \lambda_{U \to L}) \times t\right)$
2	$L \to L$	$\exp\left[-(\lambda_{L \to U} + \lambda_{L \to R}) \times t\right)$
3	$R \to R$	$\exp\left[-(\lambda_{R \to U} + \lambda_{R \to L}) \times t\right)$
4	$U \to R$	
5	$U \to L$	
6	$R \to U$	see eq. 36, if $\lambda_{A \to B} + \lambda_{A \to C} \neq \lambda_{B \to A} + \lambda_{B \to C}$
7	$L \to U$	see eq. 38, if $\lambda_{A \to B} + \lambda_{A \to C} = \lambda_{B \to A} + \lambda_{B \to C}$
8	$L \to R$	
9	$R \to L$	

The complete likelihood is the product over individuals and time points. As a result of these nine possible transitions, we have six rate parameters for each video, grouped by treatment group. We code this model in the Stan language and fit it using Hamiltonian Monte Carlo as implemented in Stan (Stan Development Team, 2014).

4.3.4 Results

In all the following results, the left hand side (L) is the side with the larger sunfish (if there were two sunfish in the tank) or the side with a sunfish (if there was only one sunfish). The sides were originally randomized but are rearranged here to ease interpretation.

One way to measure preference in this system would be to look at the symmetry of estimated transition rates within each treatment (figure 16). If shiners had no preference for one or the other side of a tank, the difference between transition rates U $\to$ R and U $\to$ L, or R $\to$ U and L $\to$ L should be about zero. For simplicity I present only these two pairs of transition rates. For all treatments except the medium-vs-short treatment the forward bias (transition out of the uncertain state, $\lambda_{U \to R} - \lambda_{U \to L}$) was estimated to be near zero, with 95% credible intervals covering zero. For the medium-vs-short treatment, there was bias towards the larger opercular flap. The picture is similar for reverse bias (transitions out of the left/right side towards the center, $\lambda_{R \to U} - \lambda_{L \to U}$) with all except the medium-vs-short treatment and the long-vs-short treatment having estimates of near zero bias. For the two treatments with non-zero bias it is estimated positive which means transitions from the fish with the larger opercular flap are more comment than transitions away from fish with a smaller opercular flap.

An alternative method of comparing treatments is to look at a given rate and check whether it differs among treatments (figure 17). The benefit of this approach is that it does not assume any given pattern of rates in the absence of a preference. For two rates ($\lambda_{U \to R}$ and $\lambda_{U \to L}$) and all treatments (figure 6fig:shiner-lambdas) the estimated rates and their credible intervals ranged from ≈ 0.025 events per minute to ≈ 0.15 events per minute. The lowest rate (≈ 0.025) was $\lambda_{U \to R}$ for the medium-vs-short treatment, which was much lower than $\lambda_{U \to L}$ for that treatment (≈ 0.15). All other treatments had approximately equal rate estimates (≈ 0.15) with broadly overlapping credible intervals. Only the short-versus-empty treatment differed, with both rates ≈ 0.075. These rates show clear differences among treatments but they do not follow any clear pattern with regards to the intent of the

treatment and associated hypotheses.

4.3.5   Evaluation

There are two conceptual difficulties with this experiment for answering the particular scientific questions presented above. First, the number of events in this data set is relatively small. A few transitions occur (as defined) and most group have many non-events contributing to the likelihood, which is consistent with the relatively low transition rates presented in figure 17. A simulation study would be appropriate for sorting out the required sample sizes and event counts required for inference of group-level differences in event rates. Second, the mental model we have is of a data set where more frequent transitions towards a given side (resulting in a higher estimated rate) indicate preference. Preference might also be indicated by a transition to one side followed by fidelity to that side. In that case the model might still be effective (e.g.-detecting fewer back-transitions) but it would only be effective if it succeeded in using second-order events to correctly classify inter-group differences.

There are also two more technical issues. The first model presented here is relatively rigid. Our data is structured by experiment and group whereas parameters were estimated by group only. As can be inferred either from the traces shown (figure 14) or the experiment-wise correlations among shiner locations (figure 15), experiments differed widely in group behavior. A hierarchical model of event rates would allow us to distinguish more easily of among-experiment variation is occluding a clearer story. An alternative model-building approach would be to estimate experiment-level parameters first and only then proceed to construct models of group level effects.

Further work on this model will include a simulation study and the hierarchical alternative, but it is beyond the scope of this section. We have demonstrated with relatively challenging data that an approximate multi-process event model can be constructed and the parameters estimated using available software with the code presented in appendix B. For one treatment we estimate a clearly asymmetric set of transitions from the initial state with a clear preference towards the longer-flapped sunfish.

4.4   Emigration event example

We illustrate a more complex event rate model with the Atlantic salmon emigration data set. The goal of this example is to illustrate how time-varying rate functions can be applied even with competing risks. To simplify the calculations we only consider an initial "alive" state A, a transition to an absorbing "dead" state, D, and a competing absorbing state for individuals who emigrate, E, prior to their death.

A simpler model might model probabilities of transitions over discrete time intervals. A discrete time model would need to approximate continuous time using small discrete time intervals to apply to antenna data, which is collected continuously. A discrete time model could also consider only piecewise constant rate functions but since continuously collected antenna data generates arbitrary intervals in time, coherent piecewise constant rate functions are difficult to construct.

In the part of the field study relevant to this section, individual fish, indexed below using i, were captured at field occasions typically spanning two weeks. The first capture date of each individual was recorded as $t_{i,1}$ and individuals were implanted with a PIT tag and released into the stream. At subsequent field occasions individuals could be identified by a code stored in their PIT tag and further recapture dates of all individuals were recorded as $t_{i,j}$.

In addition to being detectable when the fish were in hand, PIT tags can also be detected at a short distance. Two antennas placed at the bottom of the field study site detected individuals at the study boundary and any individuals nearing these antennas had their PIT tag code read and recorded along with a time stamp generating further timepoints $t_{i,j}$.

Conditioning on first capture, our model considers the time until two competing events: emigration and death. Emigration is fully observed in part of the data set when antennas were active and not observed otherwise. Death is never fully observed, but the alive state can be observed (when individuals are recaptured) and the probability of surviving without recapture bounds the amount of time an individual can be assumed to remain alive despite failure to recapture during otherwise successful electrofishing occasions.

Each transition is modeled using an inhomogeneous Poisson process with its own rate function ($\lambda_D(t)$ and $\lambda_E(t)$) whereas survival is the lack of a transition. When an individual is alive it is at risk for transition to either absorbing state and these transitions are *a priori* independent, although their estimate rate functions might be correlated.

4.4.1 Exploratory analysis

The literature on Atlantic salmon suggests that the timing of movement can be driven by seasonal environmental changes or some other seasonal signal (e.g.-(Rimmer *et al.*, 1984; Saunders & Gee, 1964)). A simple way of visualizing such a pattern in our data is to take the day of the year that each emigration event occurred and plot it by year (figure 18). Assuming a constant risk pool would suggest that emigration is heaviest in late October through mid December, a dearth of emigration in late winter, another wave of emigration starting mid April followed by a lull in emigration in the summer months. In addition to this "slow" seasonal pattern it is also clear that emigration in each year is clumped on the timescale of days—a pattern that is apparent in most years but especially visible in the autumn of 2002, 2003, and 2004. This "fast" pattern suggests that whatever signal mediates the slow seasonal signal (e.g., day length) operates on a much finer timescale.

Figure 18 is ambiguous because it does not illustrate the risk pool, which makes it difficult to assess whether the influx of newly mobile young-of-year individuals in autumn is responsible for the autumn emigration (and potentially for more of the seasonal pattern). To assess the relationship between population structure and emigration we show a subset of 250 of emigrants (figure 19) with both their first capture and their emigration date to give a sense of the risk pool and the strength of the seasonal pattern. Broadly speaking it does not appear that autumn emigration is driven by a large autumn pool of new recruits as most lines span at least a year. However a large proportion of recruits from autumn 2003 and 2004 do appear to emigrate soon after first capture, suggesting that some important drivers of seasonal emigration vary among years. To take a closer look at the connection between emigration and age we also plot an approximate age at emigration (time since first capture) versus the seasonal timing of emigration (figure 20). Both young and old fish emigrate at all times of the year so any simple correlation between autumn emigration and population age structure is weak.

While it would be possible to further visualize the seasonal timing of emigration, we defer the more detailed analysis to the modeling step. The question we want to address is not simply whether this seasonal pattern exists, but how different the emigration is (e.g.-per 100 at-risk individuals) between the summer and autumn. That question can be addressed directly by an emigration model.

One pattern observed in this data without an explanation is the clumping of emigration

events. A plausible hypothesis for the "fast" pattern observed in the timing of emigration is that environmental conditions are driving the within-season emigration decisions. Previous work has suggested that both stream flow and water temperature could driven movement and therefore emigration (Armstrong *et al.*, 1998). To better identify this fast pattern we use *daily* discharge data (daily stream temperatures during this period are generally slowly declining and so an unlikely driver for a fast pattern) and focus only on the years 2002-2004 and only on the time period after day 250 where the fast pattern clearly occurs. In figure 21 it is clear that at least in this portion of the year emigration tends to increase with almost any storm flow. At this point it is clear that flows do trigger emigration on a fast scale to some extent and we turn to quantifying this relationship rather than exploring it further.

4.4.2 Model specification

As elsewhere we specify the model using s_j to denote observed times we condition on and t_j to denote unobserved times or observations of random variables. We discuss the model for a given individual and construct the complete likelihood for the model based on the product of individual likelihoods. The model is fully specified as follows:

$$t_{1,D}, \ldots, t_{N,D} \sim IHPP(\lambda(t)) \tag{39}$$

$$t_{1,E}, \ldots, t_{N,E} \sim IHPP(\lambda(t)) \tag{40}$$

All individuals begin in the alive (A) state and can transition to either the dead (D) or emigrated (E) state. Since the final states are absorbing, transition to one absorbing state effectively censors the transition to the other state. Below we discuss the component likelihoods which arise in our calculations. All calculations are conditioned on a previous timepoint s_1 when the individual is known to still be in the alive state.

4.4.2.1 Observed alive at time s_2

The simplest likelihood is for the no-event calculation. In this case neither observable event has occurred by time s_2 after s_1.

$$\ell(\lambda_D, \lambda_E | t_E > s_2, t_E > s_1, t_D > s_2, t_D > 2_1) =$$

$$Pr[t_D > s_2, t_E > s_2 | t_D > s_1, t_E > s_1] = Pr[t_D > s_2 | t_D > s_1] \times Pr[t_E > s_2 | t_E > s_1] =$$

$$\int_{s_2}^{\infty} p(t | t_D > s_1) dt \int_{s_2}^{\infty} p(t | t_E > s_1) dt = \exp(-\Lambda(s_1, s_2) - \Lambda(s_1, s_2)) \quad (41)$$

This calculation uses the CDF of the IHPP to calculate the integral and the fact that the two risks are independent to simplify to the final form.

4.4.2.2   Observed emigration at time s_2

This observation indicates survival up to time s_2 and emigration at time s_2 which can be stated as:

$$\ell(\lambda_D, \lambda_E | t_E = s_2, t_D > s_2, t_E > s_1, t_D > s_1) =$$

$$\int_{s_2}^{\infty} p(t_E = s_2 | t_E > s_1, \ldots) p(t_D = t | t_D > s_1, \ldots) dt =$$

$$\lambda_E(s_2) \exp(-(\Lambda_E(s_1, s_2) + \Lambda_D(s_1, s_2))) \quad (42)$$

We again rely on the CDF for the IHPP and the exponential form of the two independent components. The equivalent statement with lambdas reversed also applies for an observed death and censored emigration.

4.4.2.3   Last observed at s_1, further observations attempted

A recursive calculation can arise when an observation is obtained at s_1 showing that neither emigration nor death have occurred, followed by attempted but failed observations. In this context we have a series of time points, $s_1, \ldots, s_J$, with an associated probability of detection, p_j, for each time point. In addition we have observations $x_1, \ldots, x_J$, which are either successful (1) or unsuccessful (0) with probability p_j. When an observation is successful, the likelihood follows the no-event calculation above (since either emigration or death make the individual undetectable) with an additional detection component. When the observation is unsuccessful it is ambiguous whether the individual is alive or if the

observation failed. This results in a calculation like equation 25, except with two competing processes.

$$\ell(\lambda(t), p) = \int_{s_1}^{\infty} \prod_{j=1}^{J} (\mathbb{1}_{u_j < t} p_j)^{x_j} (1 - (\mathbb{1}_{u_j < t} p_j))^{1-x_j} \times \lambda(t) \exp(-\Lambda(s_1, t)) dt \qquad (43)$$

Due to the indicator functions this integral must be calculated in a piecewise fashion with two possible (unobservable) outcomes on each segment: death or survival without capture. We write the likelihood components below, dropping parameters for convenience, and leaving the recursion implicit.

$$\ell(\lambda_D(t_D), \lambda_E(t_E), \vec{p} | t_D > s_1, t_E > s_1, \vec{x}, \vec{s}) =$$

$$Pr[t_D < s_2, t_E > t_D | t_D > s_1, t_E > s_1] + Pr[x_2 = 0, t_D > s_2, t_E > s_2 | t_D > s_1, t_E > s_1] \times$$

$$(Pr[t_D < s_3, t_E > t_D | t_D > s_2, t_E > s_2] + Pr[x_3 = 0, t_D > s_3, t_E > s_3 | t_D > s_2, t_E > s_2] \times$$

$$\ldots) \ldots) \qquad (44)$$

The recursive calculation can be terminated either at the final attempted observation time j, which is equivalent to censoring or when $Pr[t_D > s_j, t_E > s_j]$ becomes negligible. Each component of this calculation is either the no-event likelihood or the likelihood for a single death event with no emigration event on an interval. Unfortunately I know of no solution for the relevant integral.

One simple alternative is to assume a time of death and carry out the calculation based on that assumption. To be accurate this method requires high recapture probabilities which make it unlikely that an individual would still be alive after multiple recapture attempts. Our recapture probabilities are roughly ~ 0.8 much of the time which means that after two recapture attempts the probability of remaining undetected is less than five percent. First, we assume that individuals survive one recapture occasion after their final capture.

An alternative approach to the unknown time of death calculation is to avoid the integration and sample from the distribution of the time of death. This adds N state parameters to the model (where $N \approx 7000$) which is a significant computational burden for HMC but not necessarily infeasible.

4.4.3 Results

The brute-force approach of leaving an unmodified model to sample from the distribution of time-of-death in addition to model parameters was not successful. Adaptation resulted in a very small step size and the resulting sampling clearly did not adequately explore the posterior (results not shown).

In the model with assumed time of death, the emigration rate function shows seasonal variation within each year with (typically) two peaks—one in autumn and one in spring (figure 22). In the main portion of the study where the estimates are based on the largest number of at-risk individuals there is clear variation among seasons in the emigration rates. During peak emigration times of the year, emigration rates are estimated to be high (e.g., near 2-3 per year per individual) whereas during low emigration seasons the rate falls to $\approx .6$ per year per individual). In contrast, the estimated death rate is low in the main portion of the study (figure 23, $< .1$ per year per individual).

4.4.4 Evaluation

The continuous-time method successfully delineated high emigration and low emigration times of the year. When only mark-recapture data was available (1998-2001) supplemented with limited smolt trap data, the emigration and death rate functions are confounded as evidenced by the larger posterior uncertainty in both (figures 22 and 23). Estimating apparent survival would be more appropriate for this time period.

When mark-recapture data is supplemented by antenna data (2002-2006), the method is very effective at estimating the emigration rate but it is also sensitive to prior distributions. Our current prior distribution represents the idea that both emigration and death are rare in this system at any given point in time. The instantaneous death rate is $< .1$ per year per individual. This is not a weak prior and it clearly biases estimates of death rates downwards which is visible when comparing the period 1998-2001 where death and emigration are confounded to the period 2002-2006 when emigration is measured directly and death is still measured indirectly (figure 23). Alternative priors which default to a higher death rate (≈ 0.8 per year per individual) also gave results which were difficult to interpret as they result in patterns similar to the emigration posterior (results not shown) but with much larger uncertainty. With a death rate of 1 per year per individual, 40% of individuals are

expected to survive (without an event) one year and 15% are expected to survive two years. This is roughly in line with estimated survival rates from the CJS model ($0.8^4 \approx 0.4$).

My results suggest that continuous-time mark-recapture models will be useful but that using them to get accurate estimates from mark-recapture data alone requires further work. First, evaluating appropriate weak priors for continuous rate functions will be an important step forward. This should likely be part of a more comprehensive simulation study of these methods. While continuous time methods have been the subject of simulation studies, the combination of partially observed data and continuous rate functions which are not piecewise constant is distinct enough to warrant exploration as a separate method. Second, weak priors become more confusing to visualize for functions as there is only one dimension to visualize uncertainty. Developing methods for visualizing uncertainty in rate functions would go a long way towards making continuous-time methods more accessible. Third, working fully with mark-recapture data requires integrating over the uncertainty in the time of death. We were not able to successfully do that in the present work but the results (failed sampling due to a small step size) suggests that the problem is with our parameterization rather than with the method itself. Exploring the component pieces of this method (continuous rate functions, partially observed data, prior specification, competing outcomes) individually via simulations studies should lead to a resolution of the problem.
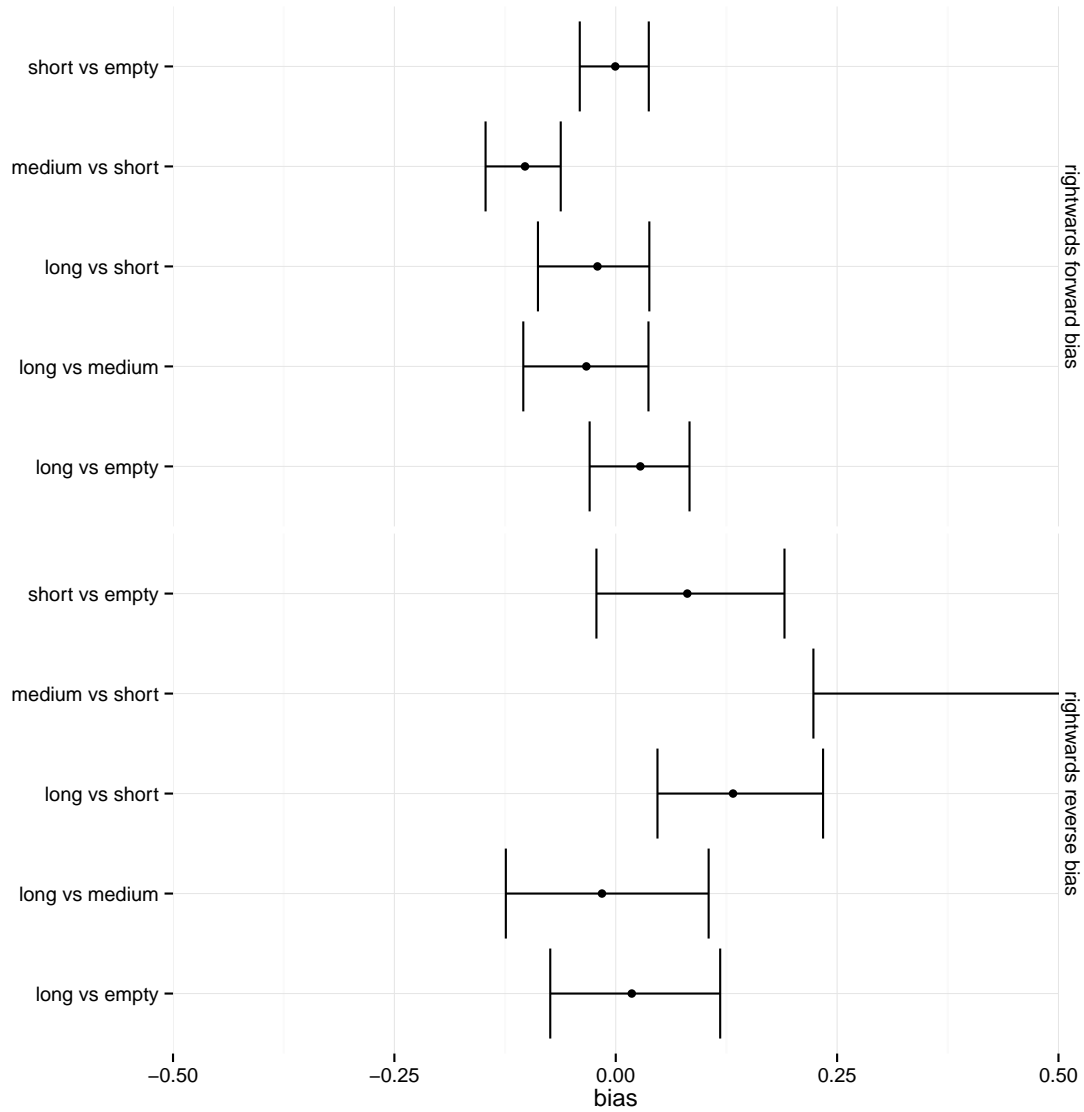
Figure 16: Transition rate asymmetry is limited. Bias is defined as the difference between center-to-left and center-to-right transition rates. Estimates are calculated directly from MCMC samples as $\lambda_{U \to R} - \lambda_{U \to L}$, and credible intervals are summarized as 2.5% and 97.5% quantiles. Bias is generally near zero and indistinguishable from zero based on the credible intervals except for the medium-vs-short and long-vs-short treatment.
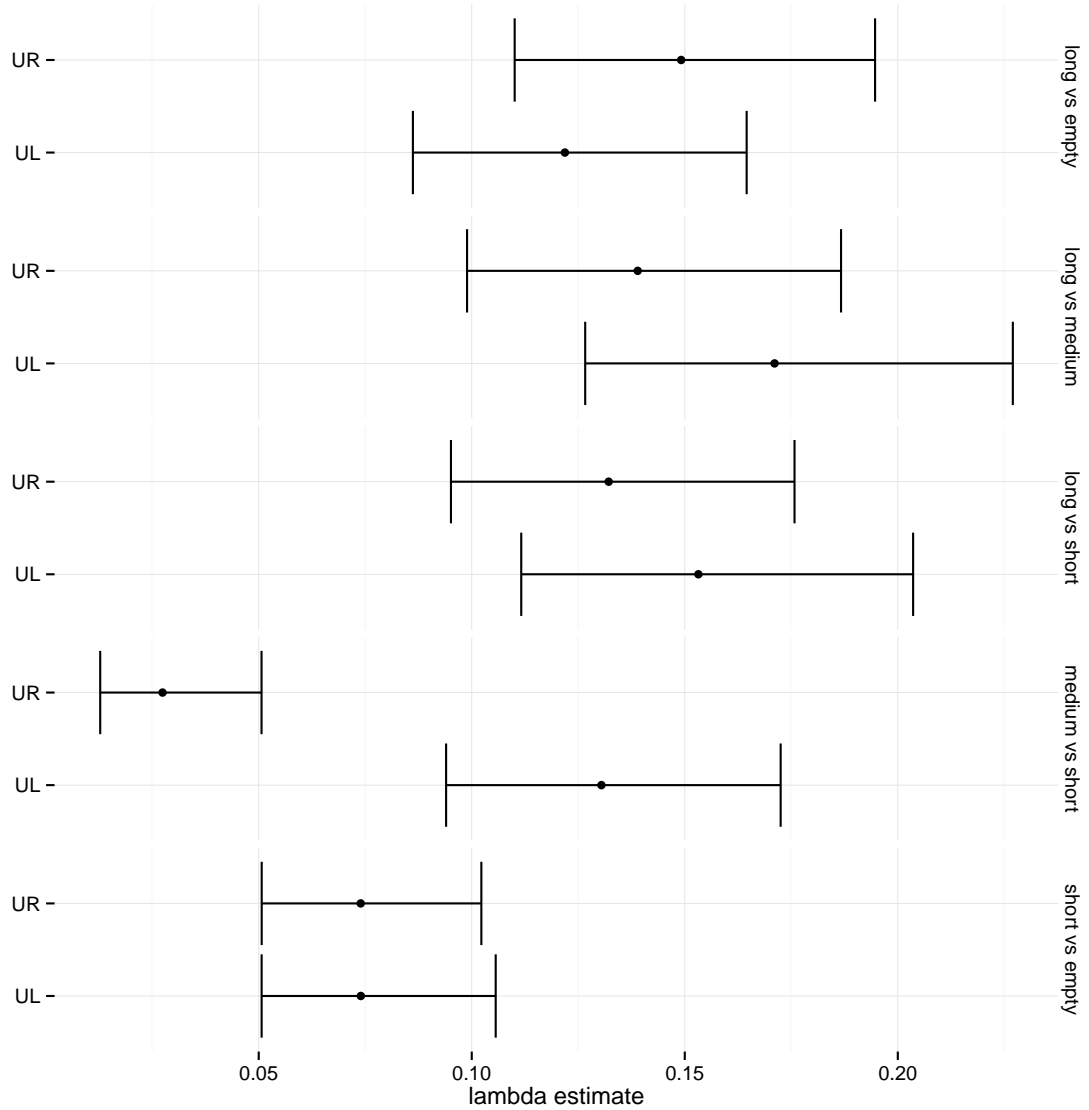
Figure 17: Shiner transition rate estimates. Estimated transition rates for each treatment showing only U → R and U → L transitions. Estimates for all two-sunfish treatments tilt towards a higher rate of transitions to the longer-flapped fish. Only one treatment showed a definitive effect (medium versus short opercular flaps) and this effect is driven by a low rate of transitions towards the shorter-flapped fish (U → R).
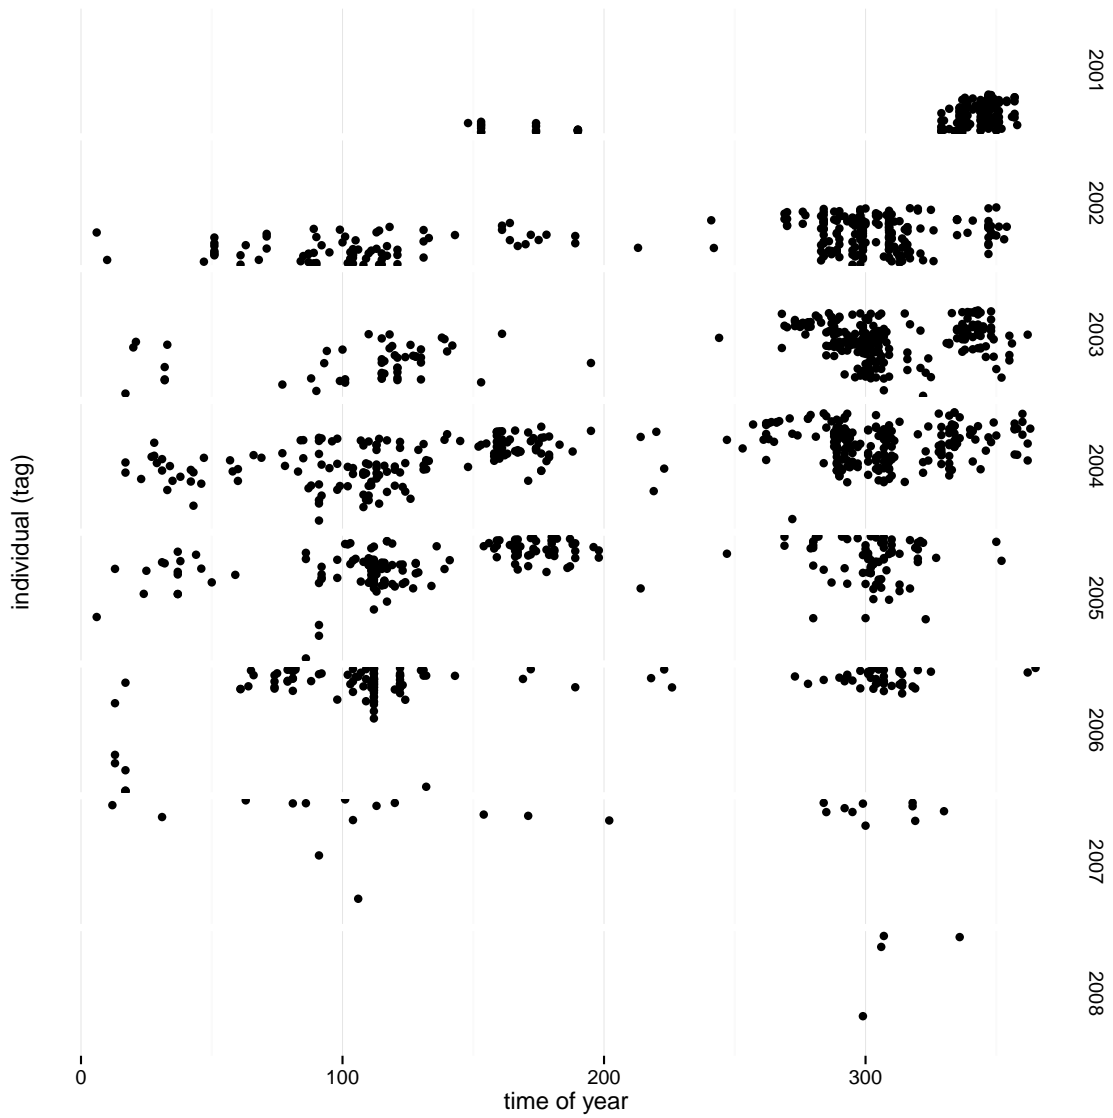
Figure 18: Antenna data shows elevated emigration in autumn. Each point represents one emigrant with the date of final boundary detection on the x axis. There is a clear pattern with late autumn emigration and early spring emigration being most pronounced. This figure does not illustrate the risk pool.

86
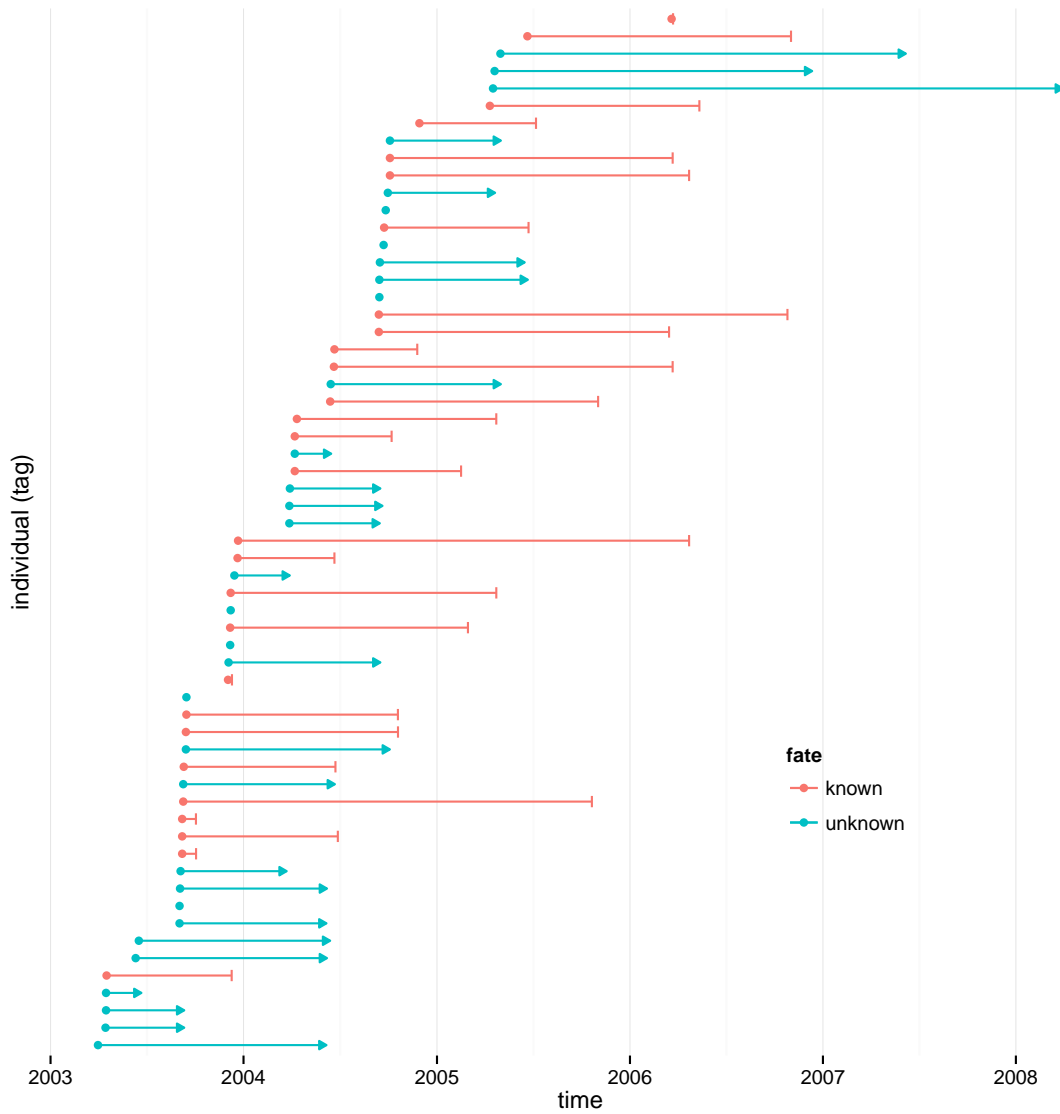
Figure 19: A random sample of fifty emigrants, their first capture times, and their final detection at the site boundary. Individuals are sorted based on their first detection time. The at-risk pool for emigration can be counted as the number of horizontal lines crossing any given vertical line and suggests that even considering the at-risk pool, there is some synchrony in emigration.
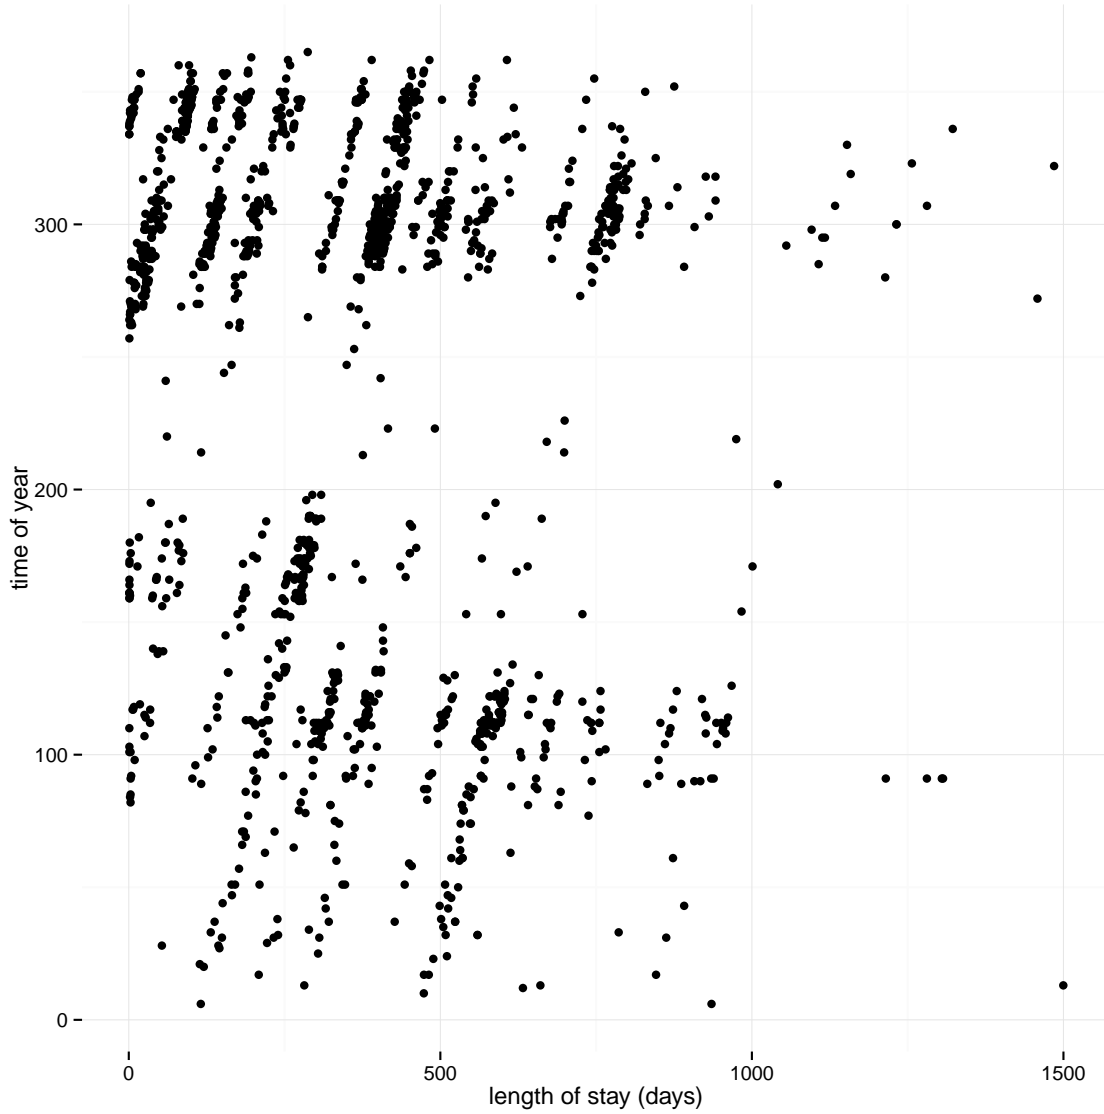
Figure 20: Age is independent of emigration time. Each point represents one emigrating individual. The x-axis is a proxy for age (time since first capture) and the y-axis is time of year. Spring and autumn emigration appear to be distributed equally over multiple age groups suggesting that age structure in autumn does not impact the strength of autumn emigration.

88
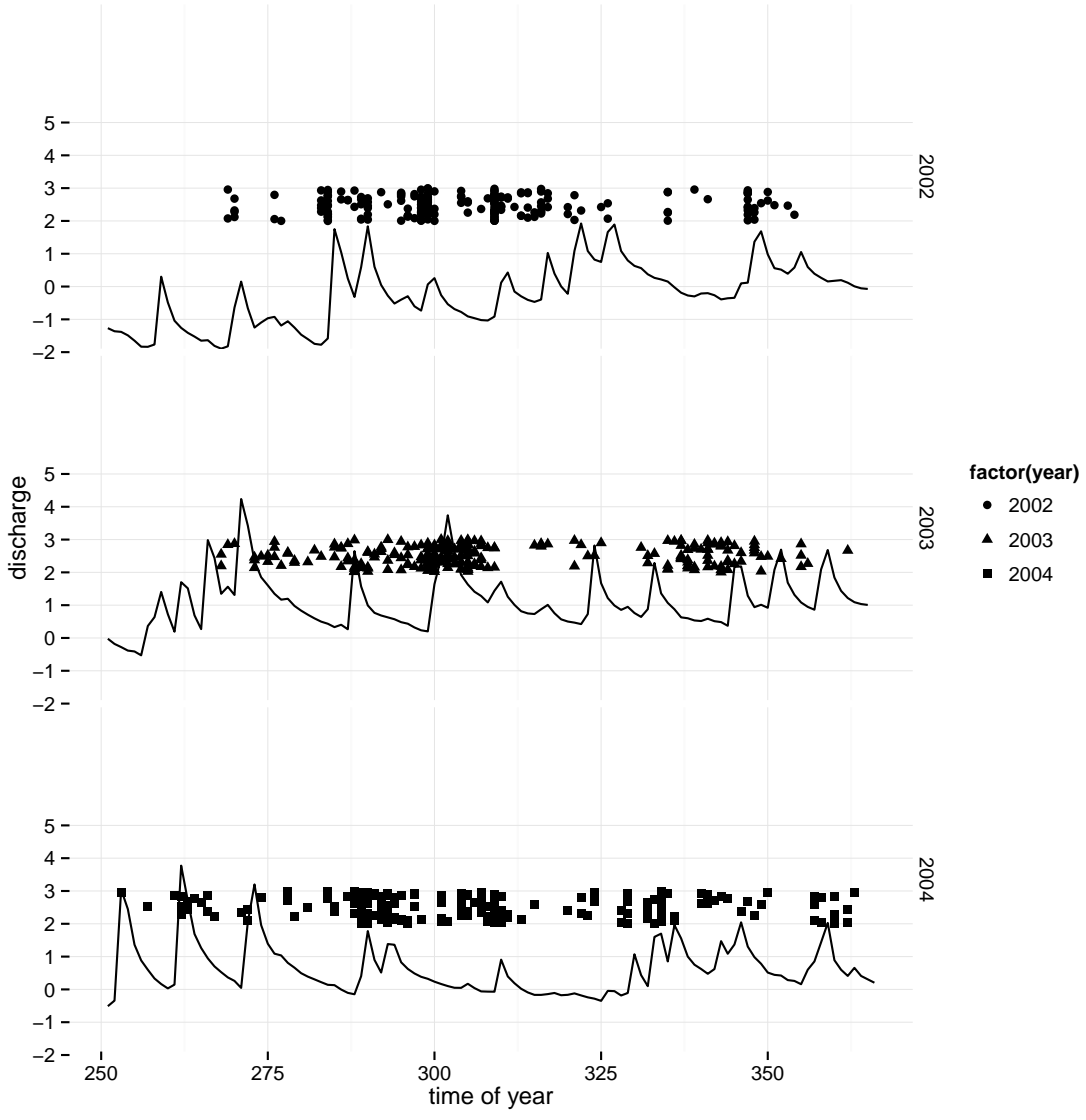
Figure 21: Emigration increases with late-season storm events. Storm events result in sharp increases in flow which then decline. The flow covariate here is a deviation from season-typical flow on the log scale which emphasizes small increases more strongly. The absolute magnitude of the peak does not appear to correlate with the amount of emigration, only the presence of an event.
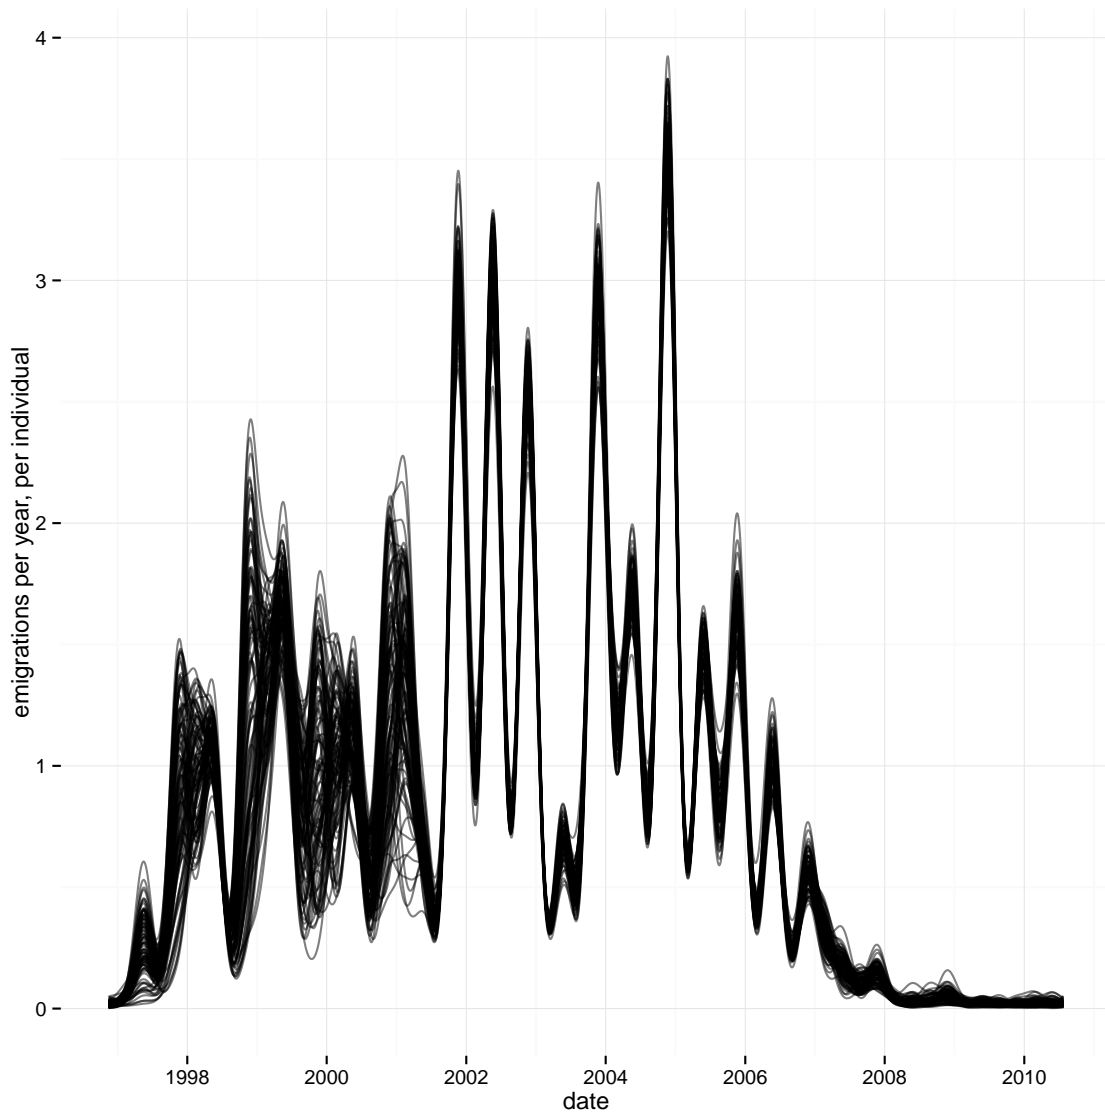
Figure 22: The estimated emigration rate function is highly seasonal. One hundred draws from the estimated emigration rate function over the course of the entire West brook study. The estimates degrade and shift towards priors at the head of the function where antennas were not active and at the tail where fewer individuals are in the risk pool.
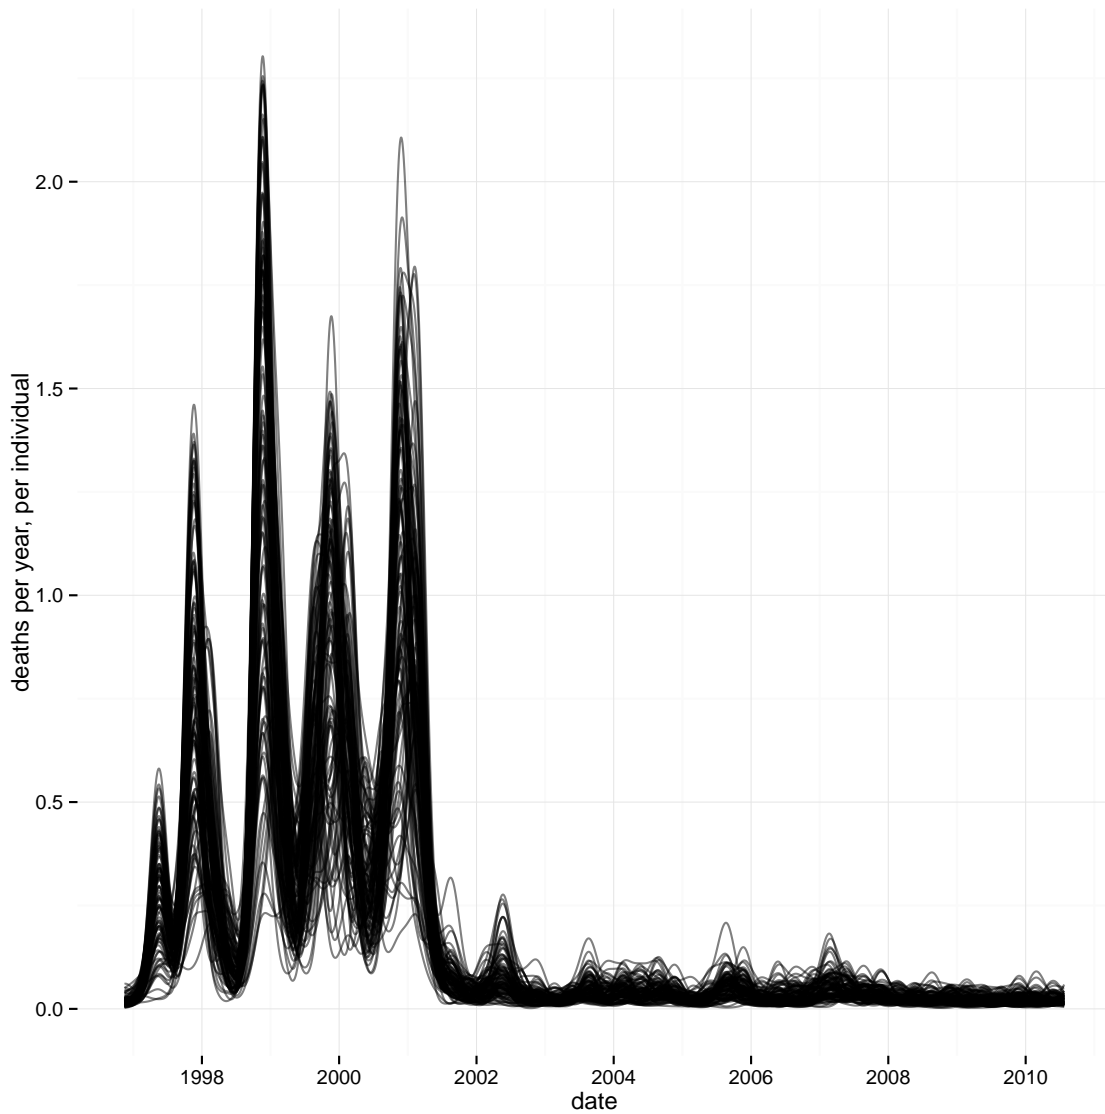
Figure 23: The estimated death rate function is confounded with emigration. One hundred draws from the estimated death rate function over the course of the entire West brook study. The estimates degrade and shift towards priors at the tail of the function where fewer individuals are in the risk pool.

CHAPTER 5

SIMPLE MODEL BUILDING FOR DEEP INTERACTIONS

Even in the case of well-studied organisms, characterizing responses to environmental co-variates can be challenging. Responses of biological systems are generally speaking not linear over the whole range of environmental conditions and laboratory studies only provide a rough guide about what sorts of non-linearities can be expected. Worse, as field studies rely on natural conditions to measure these responses, the extremes will by definition be poorly characterized. An additional complication is that many systems are seasonal and the organism's plastic response to the seasonal cycle will cause widely varying responses to environmental covariates during different parts of the year. This leaves to the analyst the task of describing typical responses, locating non-linearities, and characterizing their relevance in a context where the response may not be stationary over the course of the season.

For an important subset of models, we suggest that available methods can be applied to standardize the treatment of these difficulties and simplify the task of modeling the response of vital rates to covariates. We focus on three areas: 1) making the detection of non-linearity an estimation task rather than a model-choice task; 2) making the detection of interactions an estimation task rather than a model choice task; and 3) making the uncertainty associated with the estimated response at a point in covariate space a function of the concentration of data around that point. Our goal is to provide a method for accomplishing these tasks that is adaptable with regards to the source of the data as well as the observation process.

5.1 Local regression in biological response modeling

One of the tools we use to address this modeling problem is local regression. Local regression is a (possibly weighted) regression applied to a neighborhood in the space defined by the covariates. The benefit is simple: even in a system where responses to extreme environmental conditions is highly non-linear, the response around average conditions may be close enough to linear. A variety of other methods including Gaussian process models also have this property so the choice of the core method is less important than the local property. This section proceeds by defining the key concepts of a local regression, motivating our preference

for local regression over the construction of a more complex global model, and discussing the desired properties of confidence intervals in local regression.

Neighborhoods are familiar to ecologists from spatial models where neighboring (in the colloquial sense) locations are typically modeled as being similar either by forcing a correlation in effects on discrete neighboring location, or by forcing effects on a continuous landscape to change smoothly in space. In the present case we are interested in the space of environmental conditions where a point defines a set of conditions—for example the average temperature and humidity in a given year. A neighborhood around a point in the space of environmental covariates can be thought of as set of conditions similar to the original point. Moving from the point of average temperature and humidity to any other point in the small neighborhood around average would see conditions that are near average. We will re-use this terminology throughout the manuscript.

One difficulty with applying regression to establishing more complex relationships is that polynomial regression, the typical method for introducing nonlinearity and interactions, can produces coefficients that are difficult to interpret directly and it can easily produce misleading artifacts in prediction. This difficulty is well known in the modeling community and can be mediated by interpreting predictions rather than coefficients and avoiding extrapolation. Figure 24 demonstrates these issues with predictions from linear and quadratic models on a contrived data set.

It is possible but time consuming to avoid misinterpretation or misuse of polynomial regression models. In a simple model a trivial plotting exercise can be used to discover pathologies (figure 24). It is possible to clearly delineate how the functional form interacts with the data, whether interpretable features such as biological optima are present, and to delineate the neighborhoods in parameter space where confidence intervals are trustworthy.

In complex models with many covariates and higher-level interactions this plotting task becomes non-trivial. Two dimensional interactions suggest contour plots for visualization, which makes it more difficult to compare confidence intervals to data. If low-dimensional slices of the space are visualized instead, the choice of slices will affect whether artifacts become obvious. If the covariates space is multi-dimensional it also becomes more difficult to decide when data sparsity—the rarity of certain combinations of covariates—occurs. It simply becomes unclear whether an analyst is looking at interpolation or extrapolation. Finally, if the model relates to mark-recapture models or other models with a complex observation process, there may be no direct way of plotting the response versus the covariates.
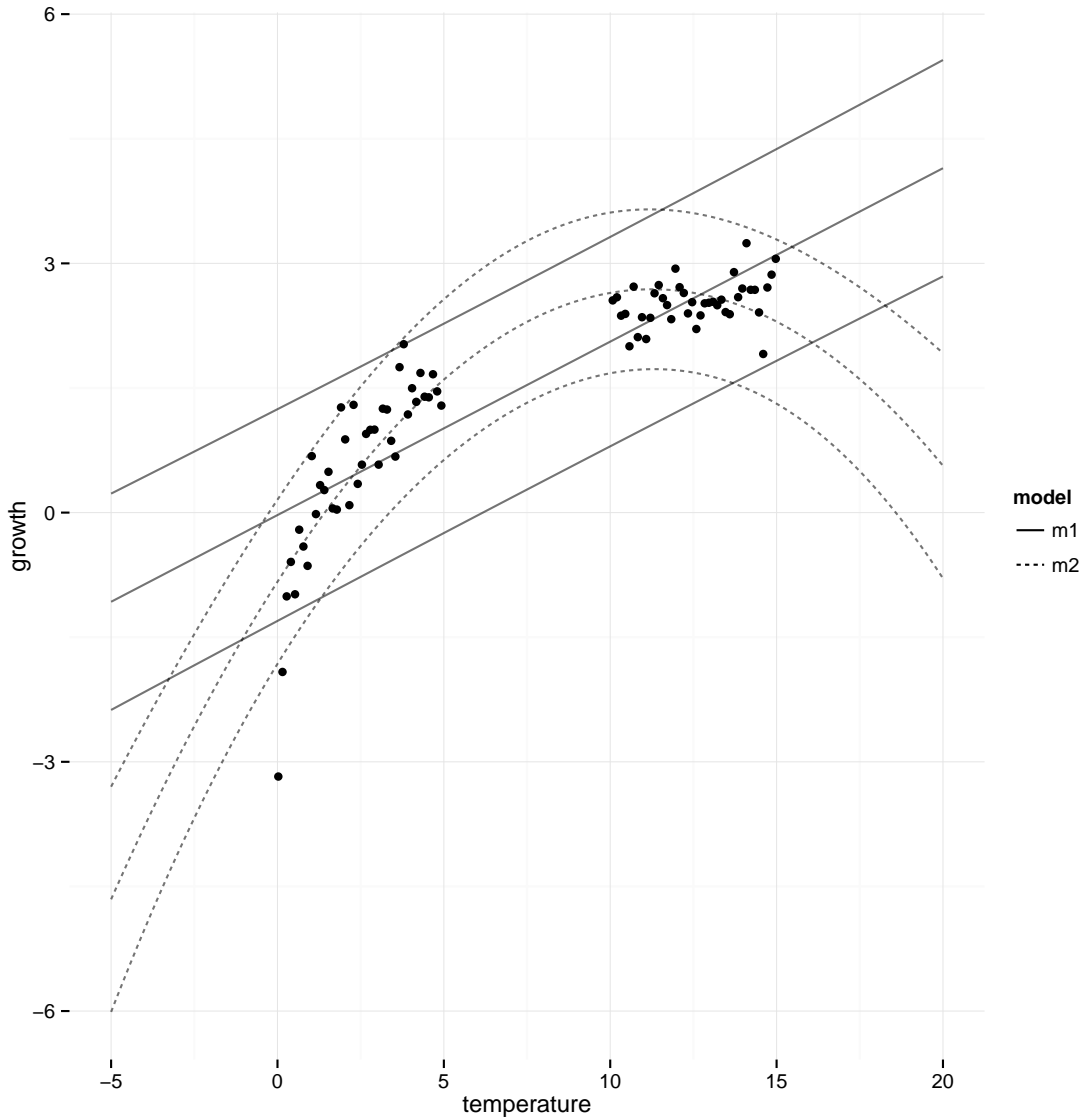
Figure 24: Prediction from linear and polynomial (quadratic) regression are fragile. Problems include: 1) in a global model, predictive confidence intervals are the same regardless of how much data is available locally; 2) the functional form must be correct across all parts of the covariate space; and 3) common interpretations of coefficients, such as a negative quadratic coefficient used to infer a maximum are very sensitive to model mis-specification.

It would be preferable to turn this model checking problem into an automated estimation problem. One desirable property would be for the model of the estimated biological response to a covariate in a neighborhood to rely primarily on data from that neighborhood. This is a core property of local regression. Another desirable property would be to have the credible intervals (and prediction uncertainty) be a function of the data available in the neighborhood of the point in covariate space for which the prediction is made. This *can* be a property of local regression given the right model structure (figure 25).

5.2   Implementation of local regression

5.2.1   Defining a covariate space for biological responses

The complications involved in estimating biological responses to environmental covariates are simplified by the restricted covariate space. For many organisms, the range of conditions that allow for persistence for even a short duration is loosely understood. This information can be gathered from physiological and behavioral studies and it is qualitative rather than quantitative, which makes it simpler to apply broadly. While biological systems do exhibit dramatic breakpoints in their responses to environmental covariates, we typically do not expect multiple dramatic reversals in response. This combination of a loosely restricted covariate space and expectations about the complexity of effects in this space simplifies modeling by restricting the expected complexity of the response function. Figure 26 displays this idea graphically for a two-dimensional covariate space.

We use a two-dimensional version of this restricted covariate space to consider the complexity introduced by seasonality. Seasonality can modify the response of an organism to a covariate. This often occurs because of behavioral and/or physiological shifts in the organism in response to dramatic shifts in the range of covariates encountered in different seasons. Often seasonality is modeled as a discrete effect and the response is estimated separately in each season. This approach is effective when the transitions between seasons are crisp *and* when the transitions themselves do not have drastic negative effects (e.g., high mortality or depletion of resources at the transition). Figure 6 shows estimated survival and its response to two covariates for Atlantic salmon in a small stream system. The response was estimated separately for each season. When seasonal changes are smooth, seasonality can be considered as a smooth circular covariate, which interacts with other covariates.
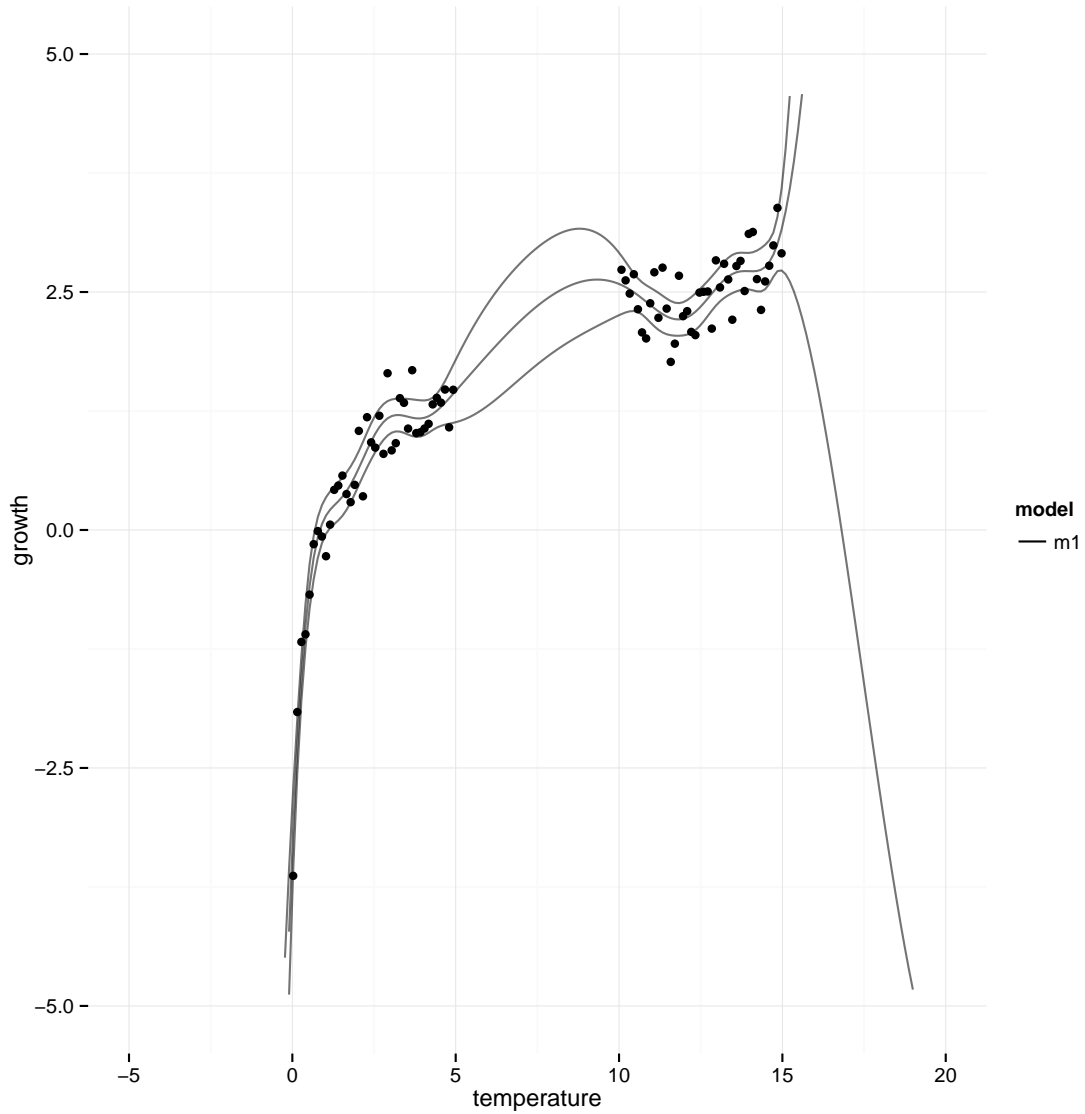
Figure 25: Local regression estimates depend on local data. Useful properties of local regression can simplify the analysis of high-dimensional data—local response is a function of local data and with the choice of the right method, confidence intervals *can* be a function of the available data. This model used the **MGCV** package with third order polynomial thin-plate splines.
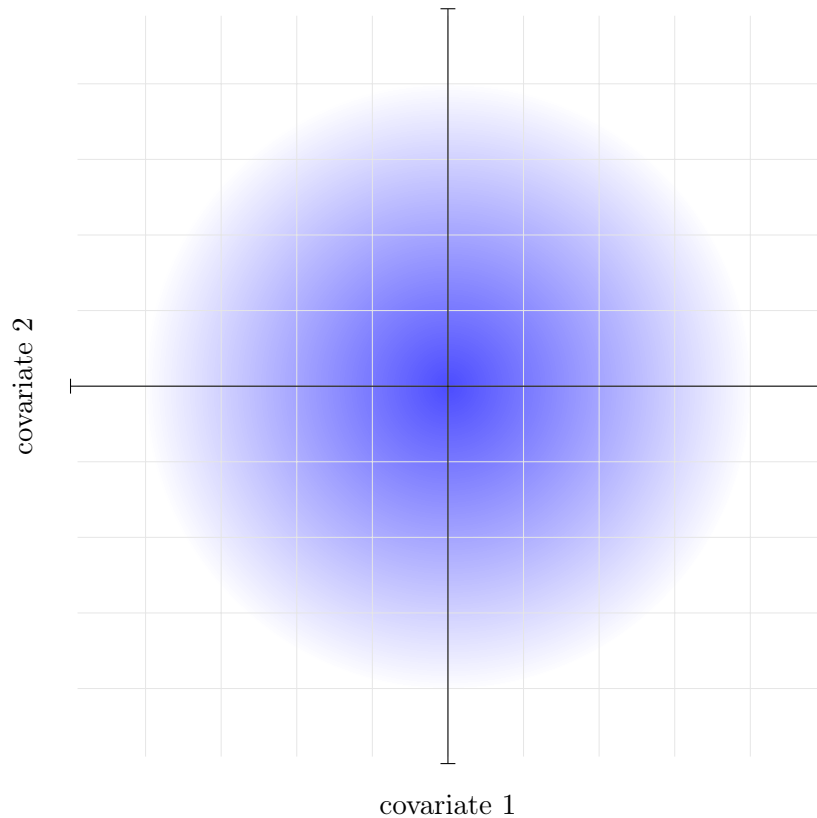
Figure 26: Biological covariates are typically restricted. The axes cross in the region of covariate space representing average conditions. The blue region displays the relevant range of conditions around the average. The response function is expected to be non-linear and interactions among covariates are also common. However, only a small number of non-linear regions and interactions are expected. The restricted covariate space of biological systems is useful for simplifying models.

The Atlantic salmon system is an interesting intermediate case because the seasons are well defined biologically, but they do change smoothly from the covariate perspective. Figure 1 shows environmental conditions as observed at the Atlantic salmon study site with a local smoother used to described (smooth) season-typical conditions and points for observed conditions. In this case the estimated responses (figure 6) for winter and summer describe

97

periods of relatively stable if challenging conditions. The estimates for spring and autumn describe periods of rapidly but smoothly changing conditions where one challenge (high water temperature and drought) is substituted for another (near-freezing temperatures and ice formation). It is interesting that both transition seasons have generally lower survival than the stable but challenging seasons. From an organismal perspective this is interesting but unsurprising since both physiological and behavioral changes are involved in the transition periods. Whether they are directly responsible for decreased survival is unknown. In any case if the transition periods are not captured any inference will be misleading.

One flexible way to treat seasonality is as an additional covariate. Using seasonality as a continuous variable lets the analyst avoid trying to partition the year into discrete seasons based on limited information. The idea of a neighborhood becomes relevant again because the effect of, for example, temperature in summer, could be thought of as an effect of temperature in the neighborhood around July 4th. Much as in spatial models, this effect of temperature in summer is linked to the effect of temperature in other locations in covariate space by correlation.

Seasonality as a smooth covariate is complicated by being circular. The effect of the first day of any given year should be highly correlated with the effect on the last day. This excludes the possibility of using standard regression models without modification but a variety of techniques for angular data are applicable. Our use of local regression simplifies modeling of circular data because any *local* piece of the season—local in the sense of a small neighborhood around a Julian day within the year—can for most purposes be treated as non-circular. The only added wrinkle is that when thinking of a neighborhood around New Year's, we must include both the end of the old year and the beginning of the new. This is easily done by first transforming Julian day to an angle and then calculating minimum angular distances as the distances between data points.

At this point we visualize our covariate space as an extension of figure 26 that shows the restricted nature of the original two-dimensional covariate space and figure 1 that shows the different parts of the covariate space covered by different seasons. The result of projecting the circle from figure 26 and projecting it along the seasonal covariate results in a bent rod that follows the covariate space explored by different seasons.

As a final twist we acknowledge the circularity of the seasonal component and connect the ends of the rod resulting in a warped solid torus shape. This space, shown in figure 27 is an example of the covariate space of interest in many modeling exercises that consider the

biological response to environmental covariates. This shape also highlights why modeling the response to environmental covariates is challenging—most predictive exercises ask us to describe the response at most points in this covariate space. Prediction exercises that consider climate change sometimes ask us to consider parts of this covariate space that are *not* explored by the climate system in any available data set.
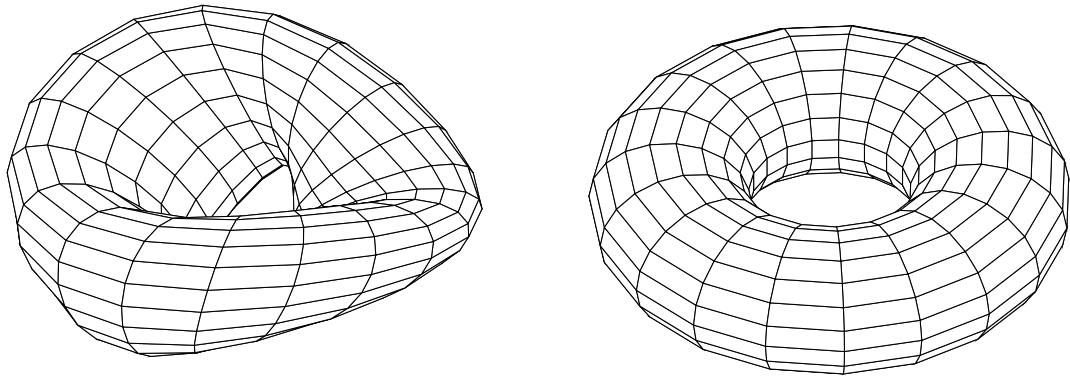
Figure 27: A solid torus visualizes the covariate space of a seasonal biological model. Average conditions and plausible extreme conditions around the average sweep out a distorted torus shape over the course of a given year. Centering environmental covariates on their typical value for a given Julian day and standardizing their variance reduces the distortion of the torus and generates a simpler covariate space.

5.2.2   Spline implementation of local regression

A generic local regression can be implemented as a spline using radial basis functions. We make this generic choice rather than choosing a specific smoother such as loess because it allows us to later elaborate the technique in response to more specific problems. Specific choices of basis functions and penalties would be equivalent to more familiar choices for local regression but we do not address those here.

In this context, a spline specifies the response function as the weighted sum of basis functions that share the same covariate space as their domain. While these functions share the same domain, they might only contribute significantly to the sum on parts of the covariate space. In other parts of the covariate space they might be effectively zero. Equation 45 shows a spline for a survival response function specified using K basis functions. The interpretation of each basis function, f_k, depends on the chosen form of the basis functions.

$$\log\left(\frac{\phi(x)}{1 - \phi(x)}\right) = \sum_{k=1}^{k=K} \beta_k f_k(x) \tag{45}$$

There is no general intuitive interpretation to the basis functions, which makes interpretation of the general spline specifications unnecessarily obtuse.

Radial basis functions provide a natural intuitive interpretation in the context of a local regression. Loosely, a radial basis function is defined as being centered on a knot point. The maximum value of the function is near the knot point and it decays away from the knot point. A Gaussian function centered on a knot, shown in equation 46 is a common radial basis function.

$$y = \exp(-\frac{1}{2}(x - x_k)^2) \tag{46}$$

This function is equal to one at x_k, positive near x_k, and approaches zero elsewhere. A gaussian function in a multidimensional covariate space is a product of the one-dimensional gaussian functions. The intuitive interpretation of a radial basis function is that it represents the response in the neighborhood around a point in a (multidimensional) space.

5.3 Evaluation of data adequacy and uncertainty

Most models are wrong and most models in biological field experiments are very wrong. Biologists typically only trust models as summaries of the data and are wary of extrapolation. In a single dimension, data adequacy for a regression is often checked with a direct comparison or rug plot. In figure 24 we see two linear regressions, each of which provides a valid summary of the data and interpretable effect sizes in a portion of covariate space. In other parts of the covariate space the implied results are clearly questionable.

Three dimensional relationships such as in figure 6 are more difficult, but an overlaid scatter plot can still present the regions of covariate space where estimates are supported by the data. If the estimates relate to directly measured data then color coding can be used to visualize the residual deviation of each data point from the estimated surface in parallel to figure 24.

In higher dimensional spaces, such as the covariate space defined for the solid torus model, direct comparisons become even more difficult but the same questions remain: 1) which parts of an estimated response surface are supported by data; and 2) if the estimates relate to directly measured data, can the deviation be visualized? Typically, results in three or more dimensions are visualized as one or more three dimensional slices drawn as heat maps or contour plots.

We also use the slicing solution and suggest addressing data-adequacy questions using kernel density estimation to produce a mask or transparency value for contour plots and heat maps. Kernel density estimates are alternatives to histograms, which generalize better to moderately high dimensional spaces. The estimates are constructed by placing a kernel function, much like our radial basis function, at each data point. Once these kernel functions are placed, the density estimate at any given point in covariate space is defined as the sum of the kernel contributions from all data points. Since the kernels are typically non-zero only in a neighborhood around each data point, the density estimate is largest where the greatest concentration of data points occurs. The multi-dimensional version is given in equation 47.

$$f(x_1, \ldots, x_K, \sigma) = \frac{1}{N} \sum_{i=1}^{N} \left[\prod_{k=1}^{k=K} \exp(-\frac{1}{2\sigma^2}(x_{i,k} - x_k)^2) \right]$$

$$= \frac{1}{N} \sum_{i=1}^{N} \left[\exp(\sum_{k=1}^{k=K} -(x_{i,k} - x_k)^2/\sigma) \right] \quad (47)$$

One tuning parameter in kernel density estimates is the band width, σ, which is chosen subjectively. When using kernel density estimates in the context of local regression, a natural choice for the band width is the scale of the radial basis functions. This assures that in regions defined as data-sparse by the estimation procedure, the kernel density estimate approaches zero whereas data-rich regions have a large kernel density estimate.

If choosing the appropriate scale parameters for a local regression is difficult, plots of kernel density estimates can be informative. For a series of increasing band width values,

a series of density plots can be used to determine when a given scale of estimation allows adjacent data points to all inform the response at a give knot point.

In a Bayesian local regression using weak priors, a natural reference for posterior uncertainty is prior uncertainty. Samples from the parameter priors propagated through the local regression machinery produce a set of *a priori* plausible response surfaces (figure 33) from the survival prior distribution. A set of plausible surfaces serves as a relevant baseline for further calculations. A posterior surface from an informative experiment should produce a set of less variable surfaces with more persistent patterns over repeated draws. These can be readily examined qualitatively by comparing small multiples of draws from each distribution.

A quantitative evaluation can proceed by comparing prior and posterior uncertainty. As an example, I use the ratio of the width of the 95% posterior and prior credible intervals. These (scaled) ratios can be visualized (figure 35) as a black mask that shades out parts of the posterior with higher relative uncertainty.

Uncertainty could also be evaluated using the width of the posterior credible intervals directly rather than in comparison to the prior intervals. This would produce a mask with reference to some absolute sense of whether an estimate is precise or not. In either case there is a reference as the analyst must make a choice of how to set the minimum and maximum color values for the mask. Referencing the prior to produce a figure focuses on whether the experiment improved our knowledge over the weak priors and referencing a desired accuracy is more informative when there are specific requirements on the use of the parameters.

Uncertainty evaluation is complementary to data-adequacy evaluation in a two-stage process. Data adequacy addresses whether any information has been collected about a particular part of parameter space. It is dependent on the covariates and, importantly, independent of the particular model used. Uncertainty evaluation is dependent on the covariates, response, priors, and the model. While this relative uncertainty evaluation is more subjective, it tells us whether the data contributed to an improvement in the state of our knowledge about a process.

5.4   Temperature model example

Our goal with this example is to illustrate the issue of deep interactions in a simple setting and how to approach it with splines. The data comprised time series of water and air tem-

peratures from eight sites in New England. Each site had a limited time series covering only part of the year. One known important predictor of water temperature is air temperature (Kanno *et al.*, 2013), although the relationship between air and water temperature can be non-linear and mediated by other covariates. We also have minimum and maximum air temperature data for each daily mean water temperature measurement. Additionally, at the site level we have a variety of covariates relating to local and upstream site characteristics such as vegetation, elevation, watershed area, and development. For simplicity—and because we use a small number of sites from the larger data set—we choose only local forest cover (FC).

5.4.1 Model specification

Due to the diversity of sites present in our dataset and the abundance of short (sub-year) time series, we choose to model the rate of temperature change rather than absolute temperature. This lets us focus our statistical model on a measure of the heat-exchange process rather than a distal result. To implement this model we write the expected value of the change as:

$$\log \frac{y_t}{y_{t-1}} = s(S, \text{TD}, \text{FC}) \tag{48}$$

Where s is the spline function incorporating interacting season (S), temperature difference (TD) terms, and forest cover (FC) terms. The temperature difference is calculated as an approximation to the heat flux from air to water. We use the average of the daily minimum and maximum temperature and subtract the previous day's water temperature to get an estimated temperature difference for daily data. This equation can be rewritten as a regression with an offset term that simplifies model construction:

$$\log y_t = B(S, \text{TD}, \text{FC}) + \log y_{t-1} \tag{49}$$

The seasonal covariate is calculated as a transformed Julian day $[0 - 1)$. There is no need to deal with the circularity of seasonal time as our data only covers streams which freeze (and therefore lack data) in winter. The forest cover (FC) covariate is used directly. We compare four simple models as outlined in table 6. The first model aims to explore the (assumed) independent contributions of our three covariates, which may be ambitious since we have only a dozen sites and time series for only part of the year. The second model

103

Table 6: Stream temperature model specifications. All models also include an offset for the previous day's log mean temperature.

Model #	components
1	$s(S) + s(TD) + s(FC)$
2	$s(S) + s(FC,TD)$
4	$s(S,TD)$
3	$\text{site} + s(S,TD)_{\text{site}}$

maintains seasonality as a separate term but allows us to look at the interaction between site-specific forest cover and the temperature difference term. The third model ignores site covariates and models only the daily measurements we have available—time of year and temperature difference. Finally, the fourth model looks at whether the global pattern of interacting seasonal and temperature effects is relevant across all sites or more likely to be driven by site-specific covariates.

These models are not definitive and our goal is not to generalize about stream temperature relationships. Our data do allow us to explore the relationship between linearity and hierarchical data, as well as multiple ways of visualizing this data.

5.4.2 Results and evaluation

For our first model each of seasonality, temperature difference, and site-specific forest cover we use smooth components to model each effect. Seasonality has a clear effect with persistent warming in the early part of the year (figure 28, top panel). On its own the temperature difference term has a clear positive effect on the change of temperature with a comforting zero effect right around zero difference between air and water temperature (figure 28, center panel). Finally, the forest cover effect varies dramatically over the span of the covariate, perhaps driven by the fact that we have a limited number of sites available (figure 28, bottom panel)..

In our second model (figure 29), seasonality retains its pattern of early-season warming but the interaction between the temperature-difference term appears to vary by site. All the sites show decreasing water temperatures when air temperature is below water temperature,

104

and all but two sites show the correlation across the entire range. However, two sites show a lack of correlation between air and water temperature when the air temperature is significantly warmer than the water temperature. We are not able to assess which factors drive this variation, though it is unlikely to be the forest cover covariate directly as neither site has the larger forest cover covariate.

Our third model (not shown) demonstrates that this global pattern, while interesting, may not be hold at all sites. The site-specific interaction model shows a remarkable number of sites have a relatively linear interaction between seasonality and temperature difference terms, with some showing only TD-based patterns and no effect of seasonality.

Our final model (figure 30) ignores site variation in order to construct a pattern relevant across sites. We see that early in the season temperatures are expected to increase unless the air is around five degrees cooler than the mean water temperature. As the season goes on this buffering disappears and the temperature difference acts in a more linear way until the end of the season, with some bias towards cooling even when air and water temperatures are closely correlated.

When analyzing a larger set of sites, the benefits of combining information across sites increase relative to this small dataset. With that expectation in mind we return to the global season-by-temperature model and apply our local regression method to it.

5.5   Survival model example

The goals of this model follow the goals of the survival model from chapter 3 where we focused on estimating age-specific and seasonal environmental effects on survival. As age structure had only very limited effects, here we group all ages together and only consider seasonal environmental effects. This model could be elaborated to consider age structure but we do not extend it here.

5.5.1   Model specification

The general implementation of the survival model follows section 3.1. The two modifications we make are that we replace the linear model for survival by a local regression spline using radial basis functions and use standard normal priors.

The spline specification must be modified to apply to a survival model. The radial spline specification gives a surface of responses and it is straightforward to evaluate at a point on the surface. The equation for this surface in terms of temperature (t), flow (f), and seasonal time (s).

$$f(t_i, f_i, s_i, \mu, \sigma) = \sum_{k=1}^{K} \left[\exp\left(\frac{-(t_i - \mu_{k,t})^2}{\sigma_t^2} \right) \times \exp\left(\frac{-(f_i - \mu_{k,f})^2}{\sigma_f^2} \right) \times \exp\left(\frac{-(s_i - \mu_{k,s})^2}{\sigma_s^2} \right) \right]$$

(50)

Survival data are collected over intervals rather than at a point. For the environmental covariates we deal with this by transforming the covariate to generate a well-defined covariate independent of seasonality. For seasonality, there is no covariate to transform so instead we take the average value of the surface over the interval from s_1 to s_2, as shown in equation 51.

$$\frac{1}{s_2 - s_1} \int_{s_1}^{s_2} f(t_i, f_i, s_i, \mu, \sigma) ds =$$

$$\sum_{k=1}^{K} \left[\exp\left(\frac{-(t_i - \mu_{k,t})}{\sigma_t} \right) \times \exp\left(\frac{-(f_i - \mu_{k,f})}{\sigma_f} \right) \times \right.$$

$$\left. \left[\frac{(\mathrm{CDF}_N(s_2, \mu_{k,s}, \sigma_s) - \mathrm{CDF}_N(s_1, \mu_{k,s}, \sigma_s))}{s_2 - s_1} \right] \right] \quad (51)$$

Here time is seasonal, so subtraction must correctly calculate the length of the interval. To correctly calculate the length of intervals that cross year boundaries, we make all calculations in calendar time and only afterwards collapse knots for a given part of a season in different years into one knot by summing their columns. In most computational environments the normal CDF, denoted as CDF_N, is efficient and available, so we use it to represent the required integral.

When the interval spans multiple knots, only the portion of each knot within the interval is included in the integral calculation. The total area calculated by the integral is scaled by the interval width so that when all knots have equal weight the final function is insensitive to interval width. This allows us to relate the spline to an interval of time rather than to a point.

Using standard normal distributions for the prior specification allows the main mass of the survival surface distribution to encompass all plausible survival surfaces for our systems, including survival percentages in the single digits and high nineties. Nine draws from this prior distribution (figure 33) to illustrate the space of plausible surfaces.

5.5.2 Results and evaluation

The output of our model is a local estimate of survival at any point in the discharge-by-temperature covariate space, over any given period of time. To produce plots comparable with the original linear model with interaction terms (fig. 6), we can visualize survival over the four discrete periods used in the previous model. These four periods are visualized using our data adequacy metric (figure 34), as well as using the uncertainty of the parameters relative to the prior (figure 35).
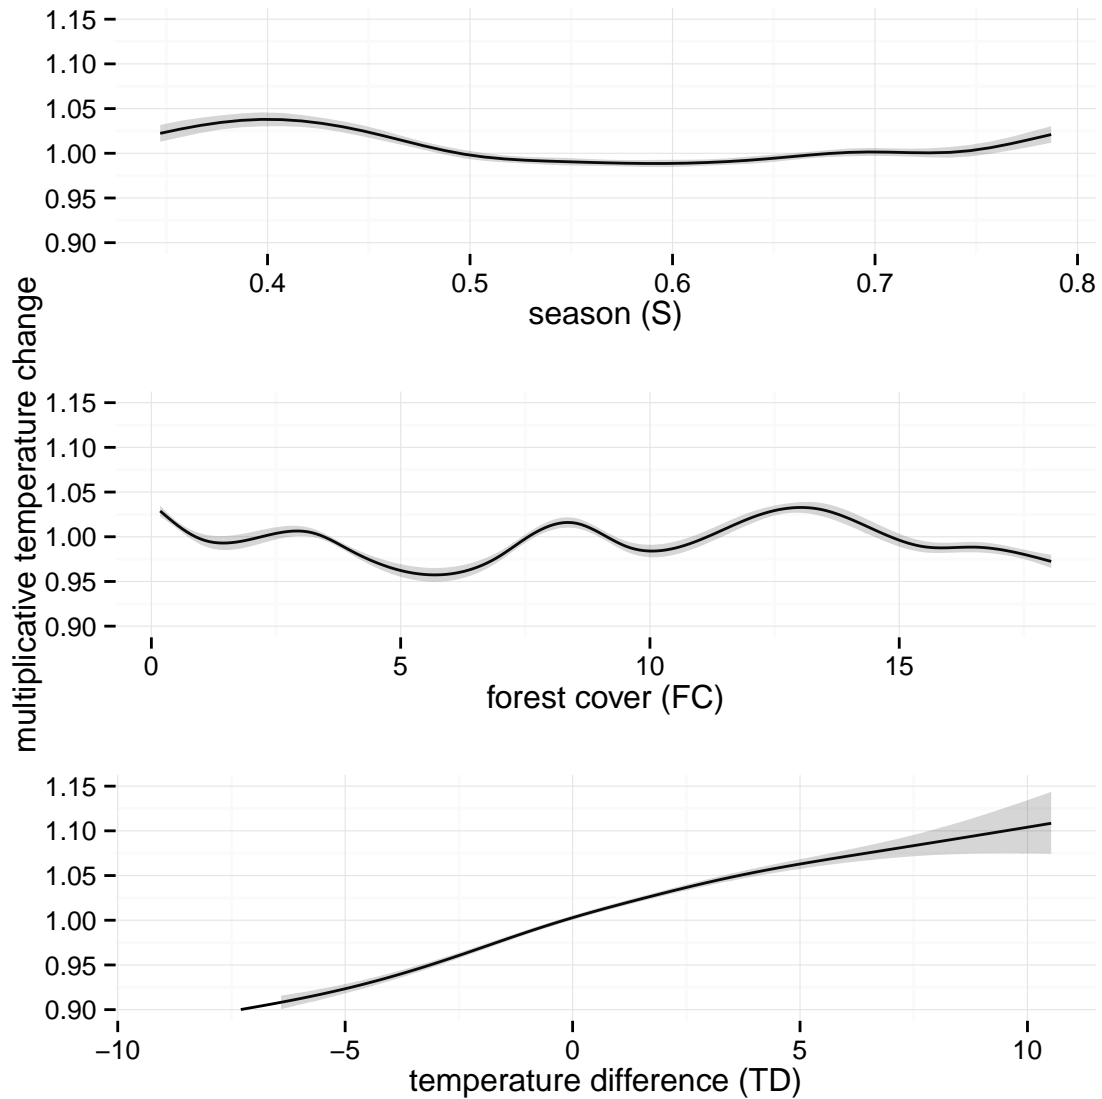
Figure 28: Effect of season (S), temperature difference (TD), and forest cover (FC) on the change in mean water temperature. The curve is a multiplicative rate of change so a value of one indicates no change in mean water temperature. The seasonal and temperature difference terms show clear patters, but the forest cover covariate relationship is clearly over-fitted despite use of MGCV for the smoothing penalty. Including a larger number of sites would solve this problem.
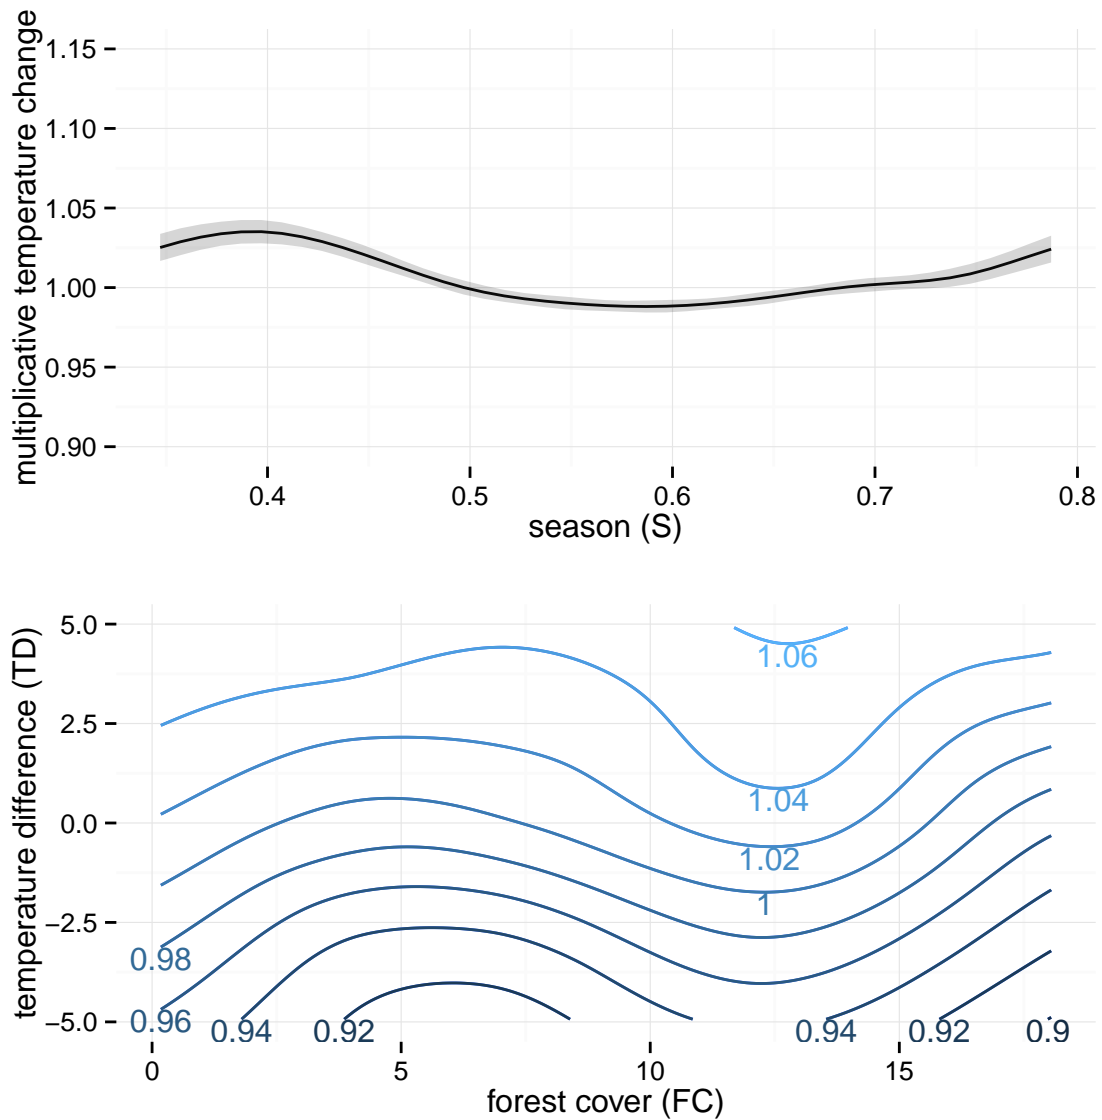
Figure 29: Effect of season (S) and temperature difference (TD) by forest cover (FC) inter-
action on change in mean water temperature. The seasonal pattern remains the same as in
the independent effects model and the temperature difference term still shows a clear posi-
tive correlation with change in water temperature. At some sites (with relatively low forest
cover values) the water temperature becomes relatively insensitive to the air temperature.
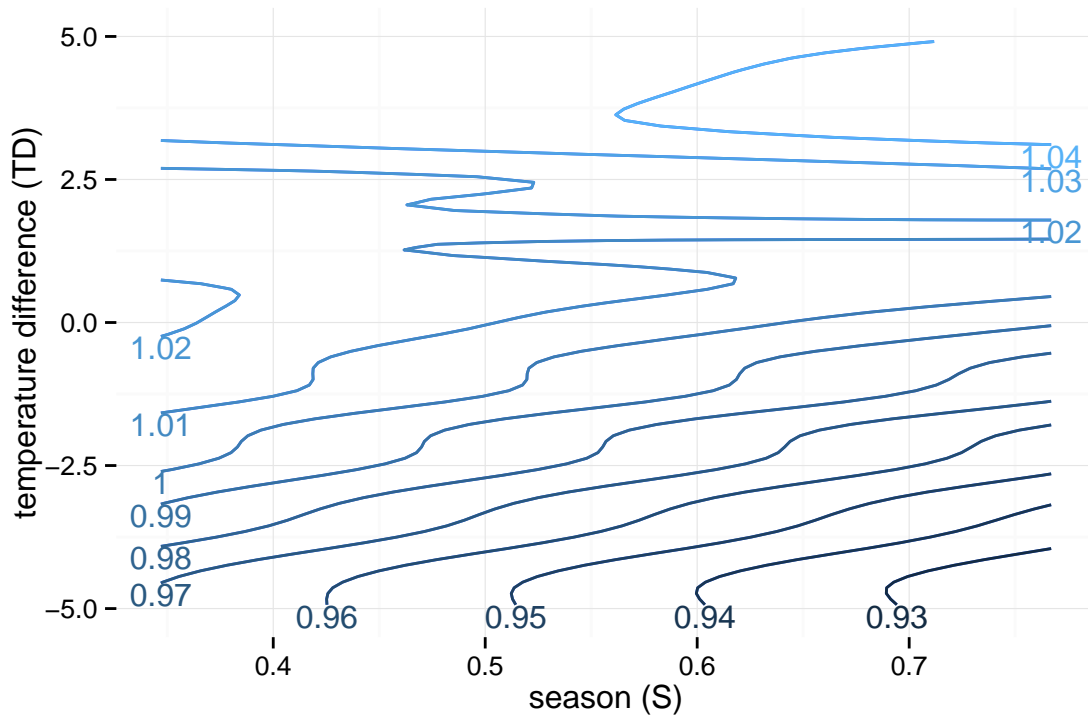
Figure 30: Effect of season (S), temperature difference (TD) on the change in mean water temperature while ignoring site-level covariates. The seasonal pattern shifts the TD effect so that early in the season warming occurs unless very flow air temperature are present.
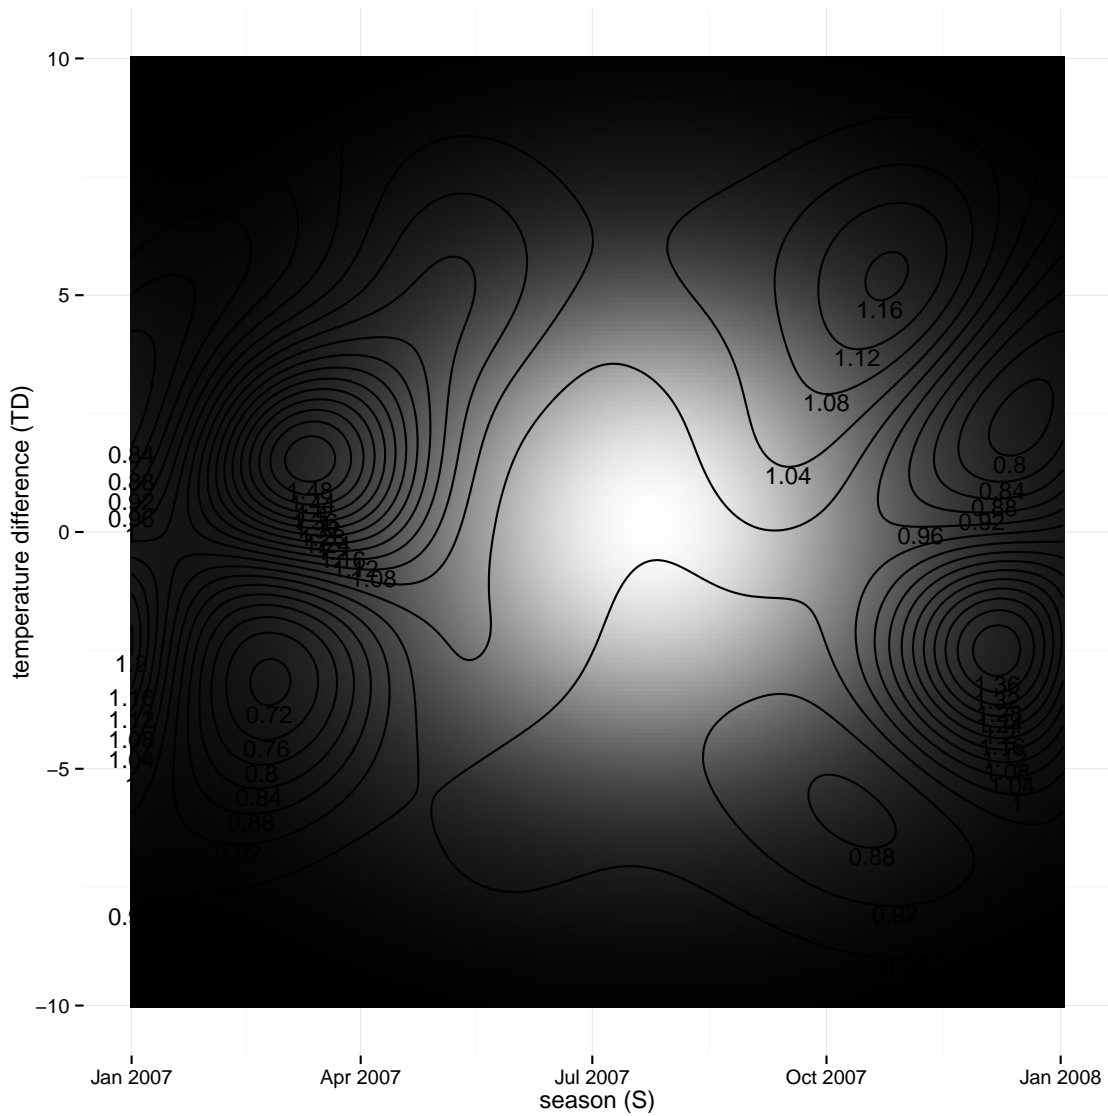
Figure 31: Effect of season (S) and temperature difference (TD) with interactions on the change in mean water temperature, with data adequacy. A large part of the covariate space (in black) was not supported by data. In the white region data was sufficient to estimate the model. More sites and multi-year time series would extend the white region to a larger portion of covariate space.
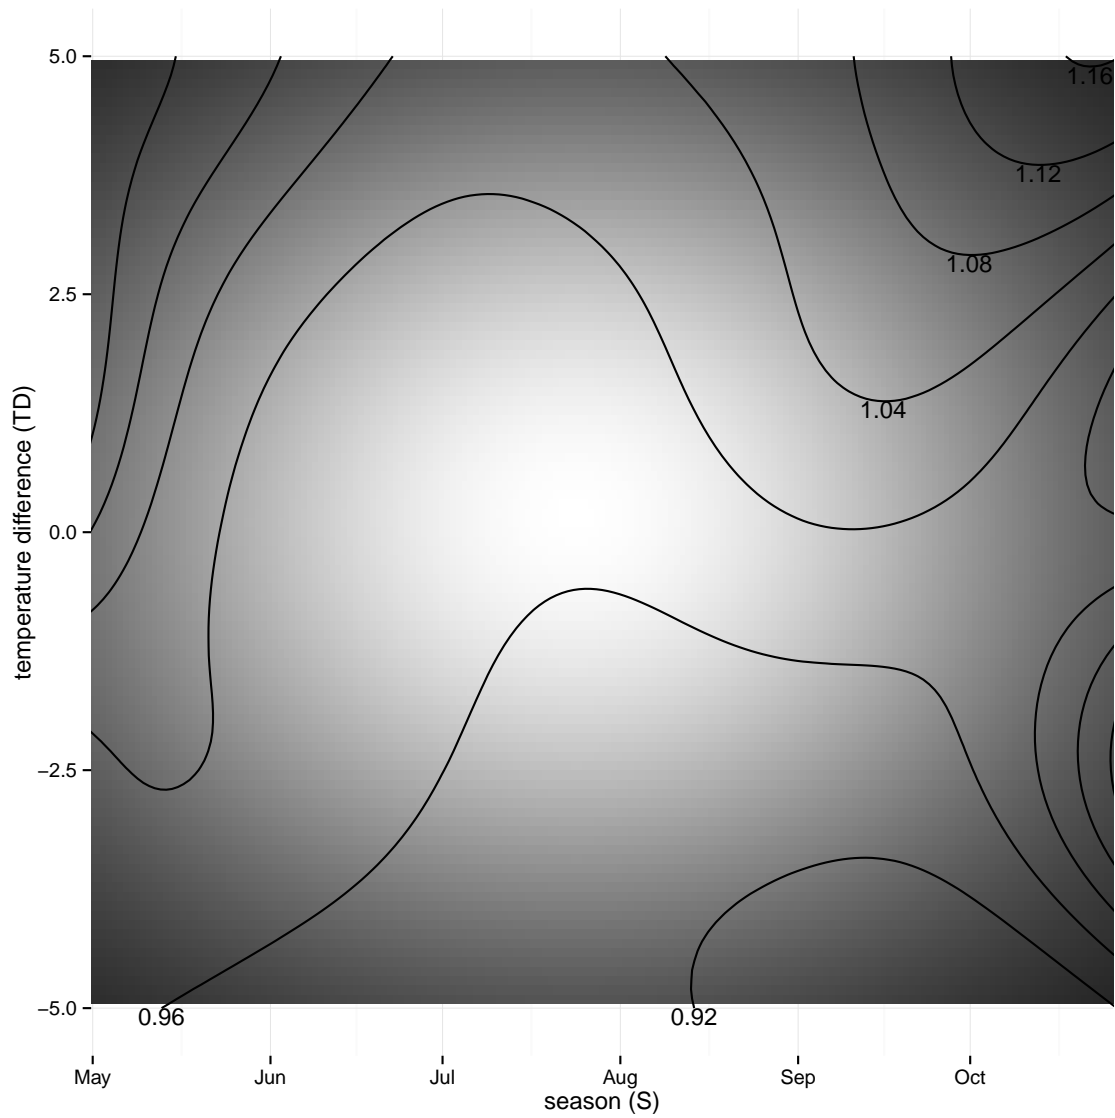
Figure 32: Zooming in on the data-dense part of the season (S) by temperature difference (TD) model. This view shows a pattern that matches closely the results from MGCV. There is a positive early season effect on water temperature rise that fades out and leaves a mostly linear temperature difference term.
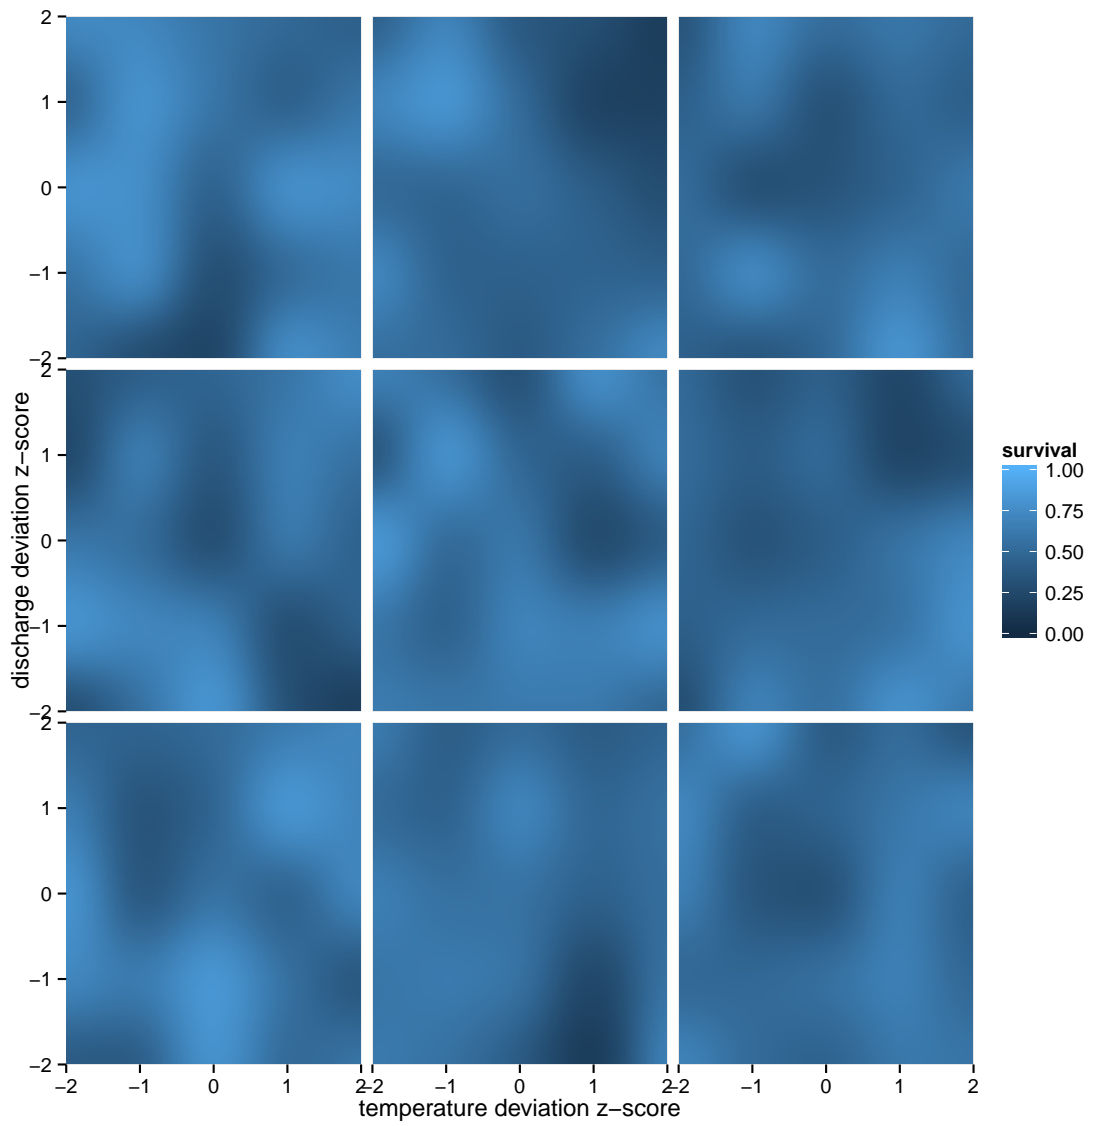
Figure 33: Nine draws from the local regression prior for the survival surface. The mean surface is 0.5 on the space shown here.
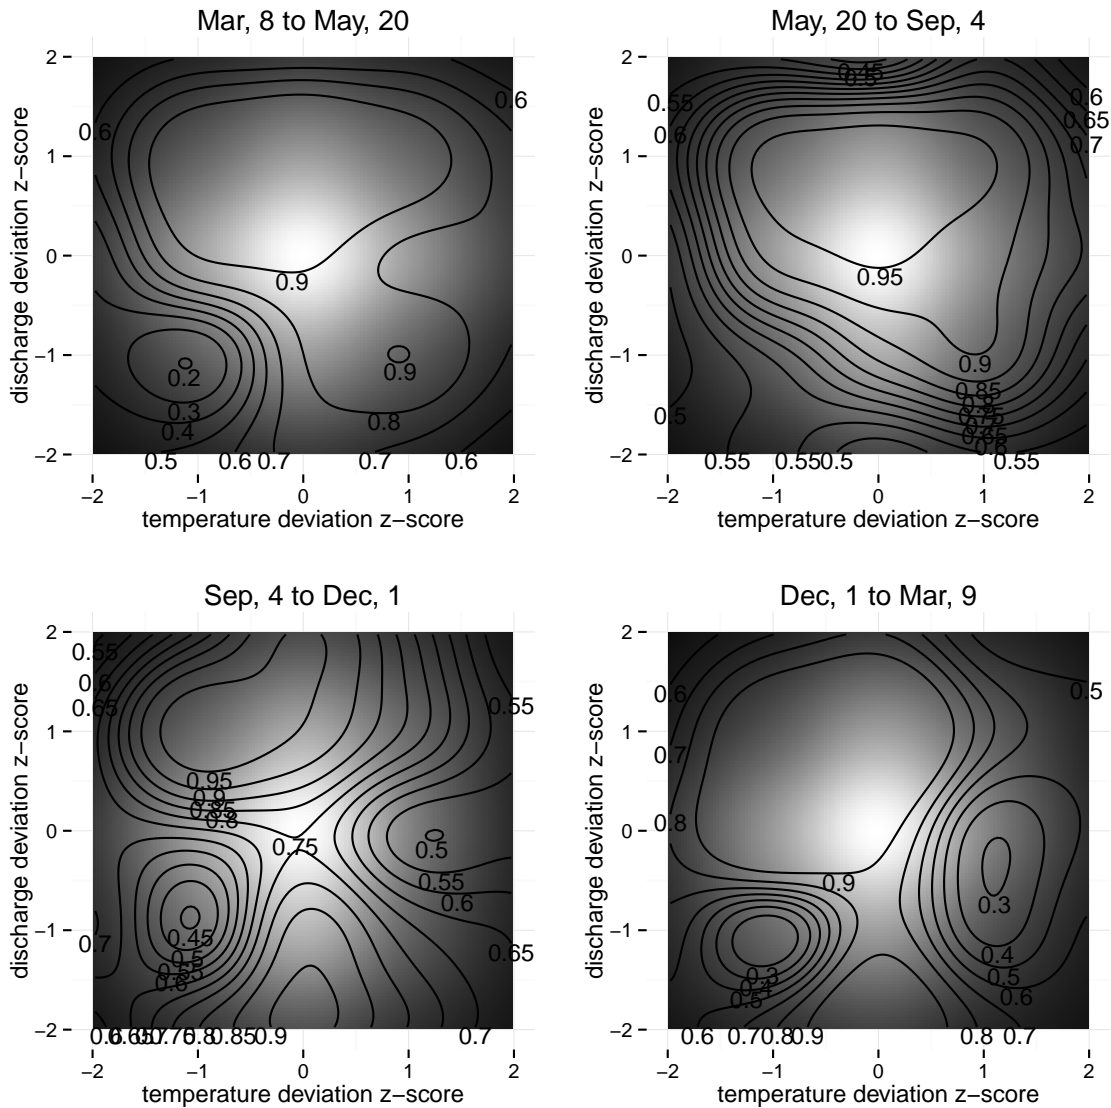
Figure 34: Effect of flow and temperature on survival for Atlantic salmon in a small stream system. Contours are labeled with estimated survival probability at the isoline. Flow and temperature show clear interactions in determining the ultimate survival probability. Unlike figure 6, the estimates here are local, less precise, but also not strongly affected by the global surface. The black mask is calculated based on the amount of data available to estimate knots in a given portion of covariate space with black regions indicating a paucity of data.
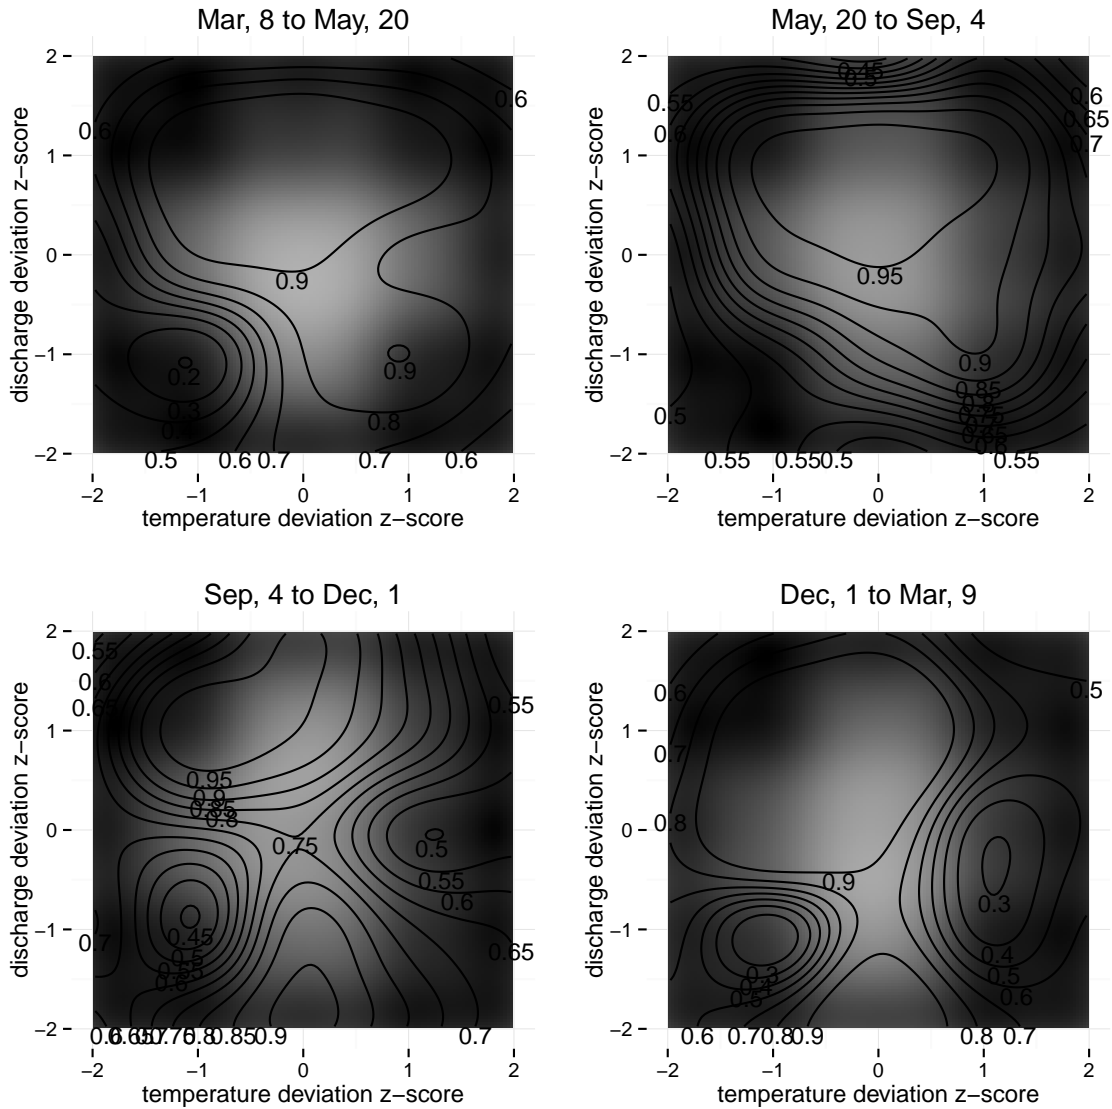
Figure 35: Effect of flow and temperature on survival as in figure 34. Contours are labeled with estimated survival probability at the isoline. Unlike figure 34, black mask is calculated based on the uncertainty of the surface in a given region of covariate space with black regions indicating uncertainty approaching the prior.

115

6.1 Environmental effects on Atlantic salmon survival

I constructed environmental response surfaces for Atlantic salmon (*Salmo salar*) that make relatively detailed statements about expected seasonal survival probabilities under a range of conditions. Even the mean surfaces presented—ignoring for the moment our uncertainty about the plausible range of estimates—suggest that temporal environmental variation as observed at the West Brook is critical to making any predictions about the population trajectory of Atlantic salmon. This result is fully in line with the findings of Crone *et al.* (2013) who demonstrate that even with a good understanding of stage-specific vital rates predictive models perform poorly largely due to environmental effects on vital rates.

In the context of the West Brook the observed environmental sensitivity is striking because it is a relatively well-buffered stream: the West Brook is shaded, a feature that results in limited daily insolation and temperature fluctuation; 2) it typically receives sufficient groundwater input to maintain good connectivity even in dry conditions; and 3) it is downstream from a reservoir that buffers both flow and temperature. In this relatively stable stream we estimated survival rates from ~ 0.95 down to $\sim .6$ or lower depending on how much trust we put in estimates on the boundaries of the explored parameter space. Clearly such variation will have important consequences for smolt production and, potentially, for the population trajectory.

6.1.1 Spring

A focus on feeding is the main behavioral feature of spring for salmon. Energy reserves, depleted in winter, are restored at a high rate in the spring (Berg & Bremset, 1998). The stream loses its ice cover (and anchor ice), and it warms from low single-digits to low teens. The stream environment also experiences some of the highest consistent flows of the year during this time period. Previous studies lead us to expect that high flows should not negatively affect over-yearling salmon survival (Bacon *et al.*, 2005; Jensen & Johnsen, 1999). Additionally there is some evidence from another analysis of this data set that growth increases with high flows in spring (Davidson *et al.*, 2010), suggesting that they are not in

general very challenging. In the global analysis, spring survival was only negatively affected by low temperature and low discharge conditions (figures 6 and 7). All other observed combinations had either a weakly negative or weakly positive impact on survival. In the local analysis the same pattern is present, with the cold dry region of parameter space having low survival whereas all other parts of parameter space (figure 34 and 35) have high survival. Both analyses show that estimates in these regions are well supported by data.

Low temperature and discharge conditions suggest a continuation of winter, delayed ice melt and especially delayed snow melt, which would otherwise drive up flows. A previous growth analysis (Davidson *et al.*, 2010) demonstrated a positive relationship between flow and growth in spring, which suggests that low flow conditions (such as from a persistent winter) might delay spring feeding and growth opportunities. The negative effect we observe for survival could play out either directly from an inability to replenish lipid reserves or indirectly through sub-optimal feeding, which leads to increased exposure to predators. Davidson *et al.* (2010) also demonstrated that shelter positively influenced growth in spring, which suggests that a trade-off between sheltering and feeding likely exists. Others have shown that fish are able to tolerate surprisingly long periods of starvation, which leads us to expect a mechanism for mortality mediated by behavior. The data on hand can not be used to distinguish among these alternatives but it does suggest that direct observation of feeding behavior and predation pressure would resolve this issue. All other combinations of spring environmental conditions improve survival. These conclusions do not apply to extreme high flow and temperature spring seasons as they were rarely observed; their confidence intervals are very broad and the estimates are not strongly supported by data.

6.1.2 Summer

Based on observed summer conditions at the West Brook, the average daily stream temperatures do not approach either lethal or feeding limits observed in the laboratory (J. M. Elliott, 1991; B. Jonsson & N. Jonsson, 2009). Diurnal temperature variation on top of high average daily temperatures may exceed the feeding limit and briefly enter the lethal temperature range (though not for long enough to directly cause mortality). Based on the environmental data alone it is an open question whether temperature can act through physiology to directly drive mortality. However, heat-stress is likely to be as elastic in fish as it is in humans and standardized lab studies might understate the extent of the effect. Summer

temperatures do exceed the thermal growth optimum for Atlantic salmon (J. M. Elliott & Hurley, 1997; Forseth *et al.*, 2001), suggesting at least some scope for a direct effect.

A second key constrain in the summer is diminishing flows that reduce the availability of shelters. During extreme low-flow years, the stream turns into a series of disconnected pools. Previous work at the same site has demonstrated that shelter is a well utilized resource during the summer season (Gries & Juanes, 1998). A variety of other studies have demonstrated that sheltering is a common response in Atlantic salmon to visual sightings of predators (Dionne & Dodson, 2002) and that individuals will defend territory more after exposure to the threat of predation (Johnsson *et al.*, 2004).

In the context of these environmental constraints, we found that summer survival was relatively insensitive to high temperatures at average or above average flows. In contrast survival was sensitive to low flows but the shape of the surface differed between the local and global models. In the global model, survival was insensitive to low flows at average or below-average temperatures, but a combination of hot dry conditions resulted in a strong drop in survival (fig. 6 and 7). In the local model, survival was also negatively affected low flows but average temperature conditions show a stronger effect than very hot dry conditions (fig. 34 and 35). In practice this comparison is difficult to make because extreme low flow and high temperature conditions were rare. In interpreting the two models I note that based on the local model, the global model is likely not flexible to accommodate the observed surfaces and the estimate of a strong effect under hot dry conditions is likely too certain.

This drop could be explained either by thermal stress or additional exposure to predation during low-flow summers. While daily average temperatures at the stream-wide gauge would appear to remain below the short-term lethal limit, temperatures at the gauge may not be representative of stream-wide conditions. Subsections of the stream with less groundwater input might depart from this gauge significantly. Additionally, daily temperature deviations that can rise to $\approx 5°C$ during mid-day might approach thermal limits. Both of these effects are exacerbated by extreme low flows that can reduce the stream to a series of pools with limited thermal connectivity. Since high flows mediate the survival effects of high temperatures, we expect that the loss of thermal connectivity might allow for locally higher temperatures that will impact survival. Another possible mechanism is exposure to predation with a loss of shelter. Shelters are clearly heavily utilized and defended resources (Gries & Juanes, 1998; Johnsson *et al.*, 2004), and shrinking flows increase competition for these and reduce their total availability. Mortality might result from competition or directly

from predation. This mechanism is not as compelling in the Atlantic salmon data set as low flow conditions lead to loss of shelter regardless of temperature conditions, but survival remains high in low-flow and low-temperature summers. Surprisingly unusually cool wet summer conditions also decrease survival but whether this is a feature of the parametric model or a mechanism related to productivity is uncertain. Summer growth is generally increased at higher flows and insensitive to temperature (Davidson *et al.*, 2010), which suggests that this response may be an artifact of the parametric model.

6.1.3 Autumn

In autumn, feeding decreases in importance, as evidenced by the choice of habitats with lower water velocities (Rimmer *et al.*, 1984). Temperatures typically fall into (and through) the range optimal for growth, making it unlikely that temperature stress is a relevant mechanism in this season. Salmon undergo a habitat shift to sheltered substrate chambers, particularly towards the end of the season (Cunjak, 1988; Rimmer *et al.*, 1983), which puts a premium on the existence of and access to suitable shelters. In results from the global model (fig. 6 and 7), survival is generally lower than other seasons but relatively insensitive to environmental conditions. Only a combination of low flows and temperatures results in decreased survival— a reversal of the summer pattern suggesting shifting needs. This pattern is similar to winter, though with weaker effect sizes (and a lower overall survival). In the local model, autumn survival estimates are also lower than any other season (fig. 34 and 35). However, the pattern of environmental effects is much more complex. Generally high flows show much better survival ($\sim$ 0.95) whereas average conditions are the worst of any season ($\sim$ 0.75). At average flows, autumns that are both hotter and colder than average show drastically decreasing survival while drier than typical conditions show improved survival. This complex pattern suggests that our global model was not nearly flexible enough and we discuss only the local model.

As the data have a discrete time structure for survival, it is difficult to assess when the shift from the summer pattern to the winter pattern occurs. High autumn flows, which lead to better survival, suggest that the typical autumn increase in rainfall arrives early. Low temperatures with average rainfall conditions suggest that winter temperatures arrive early, while flows potentially have not recovered from their summer levels. Autumn rainfall is important to relieving one of the main stressors of the summer season, whereas plummet-

119

ing temperatures signal important behavioral shifts associated with winter so this effect is unsurprising but these results do not provide a basis to resolve among the potential mechanisms. It is possible, and would be interesting to explore in future work, whether a particular set of physiological and behavioral events must occur in autumn in order to result in good survival.

6.1.4 Winter

In winter, ice formation and breakup dynamics are an important part of stream ecology (Cunjak, 1988). While the stream flows at the West Brook do not decrease again after the autumn recovery, ice formation does affect stream hydrology and habitat in a number of ways: 1) moderate ice improves shelter in otherwise bare areas by providing cover (Linnansaari *et al.*, 2009); 2) extended cold spells that generate significant ice, especially anchor ice, reduce access to shelters; and 3) extended cold spells also reduce runoff and contribute to periods of low flow that compound the decrease in shelters. Salmonids don't share shelter (Harwood *et al.*, 2002), which compounds all of these effects. The expectation from Cunjak *et al.* (1998) is that cold dry conditions should dramatically reduce shelter and therefore survival.

In agreement with these expectations, our estimates show that cold dry conditions lead to a dramatic decrease in survival. The local and global models both give similar estimates for this part of covariate space. This is a particularly important part of parameter space as it arises as a result of exceptionally cold winter periods—which also lead to decreased stream flow regardless of the amount of precipitation. In more northern salmon habitat this is likely an important source of mortality (Cunjak *et al.*, 1998) and while less common at the West Brook the mechanism appear to also operate. One of the key predictions of climate change in the northeastern United States is that winter will be warmer, wetter, and with more precipitation, which suggests improved survival.

One portion of parameter space where our results depart from previous work is that warm conditions result in decreased survival, whereas Cunjak *et al.* (1998) found that high streamflow was associated with high survival regardless of temperatures. In the global model only warm wet conditions show this effect, but that model is likely too rigid. This discrepancy with Cunjak *et al.* (1998) can be resolved by geography—due to the southern location of the study site, the warm wet conditions observed at the West Brook may not

occur at other sites where the effects of winter conditions were studied. Warm wet conditions at the west brook can involve the complete loss of ice cover and may allow predators access to the stream. This is one example of how different environments modify the mechanisms that operate in particular seasons, especially when small changes in a covariate such as temperature lead to large changes in the physical environment. Direct observations of ice cover and predator activity would help resolve this uncertainty. Interestingly, we rarely observed warm dry conditions, likely due to the winter dynamics of snow cover, snow melt, and temperature at this latitude.

6.1.5 Transition effects

One interesting aspect of local model estimates (figure 35) is that all effect estimates put the seasonal average conditions near the saddle points of the effect estimate surfaces. With the exception of autumn, there is nowhere (much) to go but flat or down. We can not make the obvious local adaptationist claim because the study site was stocked with Atlantic salmon from a regional recovery program, which makes stream-specific adaptation to the West Brook unlikely. However, average seasonal conditions do not always reside on the same saddle point. The spring and summer surfaces resemble each other closely, while the winter surface appears as a less extreme version of the autumn surface.

Somewhere in seasonal time there must be at least two transitions from one regime to the other and back again. The salmon, presumably, do as well as they can to adapt behaviorally and physiologically throughout the transition and we do in fact have some information on their shifting needs—for example towards feeding in the spring and towards shelter in autumn. The temporal resolution of the study is interesting here because autumn appears to fall in one of these periods defined by shifting needs and a sub-optimal response to typical conditions. Ultimately the present data and model can only pose this question but not resolve it. We expect that future targeted data collection and modeling would be able to improve our understanding of these transitions.

6.1.6 Atlantic salmon summary

Clearly a long-term study such as this one is difficult to replicate. Other integrative efforts in the literature (Chandler & J. D. Clark, 2014; B. H. Letcher, Schueller, *et al.*, 2014; Schaub & Abadi, 2011) have focused on improving the practical efficiency of mark-recapture studies by

combining them with less informative presence-absence or count data to better characterize a variety of sites. Our ultimate goal is different: we ask whether parameters from a detailed study can be combined with a mechanistic understanding of organismic biology to make the effects we infer transferable to other study sites and other times.

Overall the effects we observe are in accordance with expectations from the known literature on how environment affects Atlantic salmon survival. Our analysis helps place these mechanisms in the context of the seasonal cycle, suggest when they may be most important, and how the relationship between these mechanisms could be resolved. Some mechanisms can be resolved solely based on the two consistently measured environmental covariates. For example, lack of summer habitat due to low flows by itself is unlikely to cause decreased survival since we only observed such an effect when low flows combined with high summer temperatures. With additional auxiliary data such as food availability, fine-scale temperature measurements, or predator activity observations the mark-recapture data set would be able to address even more of these mechanistic questions.

In other cases, such as autumn survival important questions about the key mechanisms remain unanswered. Our work does suggest that we need to better understand the sources of mortality during this season, as well as the behavioral changes and potential knock-on effects from summer stress. We do however have the benefit of pointing to a specific slice of time and question, namely: what about low autumn flows and temperatures causes increased mortality on top of low autumn survival?

The question at hand is how combining a detailed study from a single location with the known effects from the literature can be used to improve our ability to make climate change predictions on large scales. Is there information here to link environmental effects to mechanisms that could be applied range-wide rather than needing to do site-specific detailed studies.

We expect the answer to be a qualified yes. We are seeing effects that are consistent with the published literature, but some departures that can be explained by considering site-to-site differences. For example, the stream study site has significant groundwater input and relatively mild temperatures so only high-temp/low-flow events cause summer temperature stress. Our stream is in an area with relatively mild winters so only cold/dry conditions cause significant habitat loss. However, warm conditions are warm enough to perhaps allow predators access to streams, whereas in northern streams (e.g., as discussed by (Cunjak *et al.*, 1998)) do not see this effect. How less intensive short-term work could be used to

adapt these estimates to other locations is an open question partially addressed by recent work on integrative models.

6.2 Lessons for the practice of prediction

I began this dissertation by talking about the practice of prediction, its importance to ecology, and how the literature as it stands informs our understanding of best practices. One clear lessons from Crone *et al.* (2013) is that incorporating environmental effects into prediction efforts is critical, although the question of how to incorporate environmental effects is left wide open. The current literature about species distribution models in particular suggests that prediction efforts would be more robust if they relied on more direct covariates (*sensu* Guisan & Zimmermann (2000)) and made use of existing knowledge on the physiology and behavior of organisms under consideration. These general guidelines leave open the question of how to apply knowledge from physiological and behavioral studies to prediction.

One apparently straightforward avenue for implementing the suggestions of Crone *et al.* (2013) and Guisan & Zimmermann (2000) would be to rely more heavily on Individual Based Models (IBM's) for prediction. These models can make use of trait estimates available in the literature to reconstruct the response of an organism to a variety of environmental and competitive conditions as well as their general life history. A well-studied species such as Atlantic salmon ought to be well suited to this approach.

Our experience with Atlantic salmon at the West Brook leads us to believe that an IPM-based approach would fail, and likely fail spectacularly, as complex population predictions tend to do. Despite relatively mild conditions at the West Brook, our estimated response surfaces for survival alone demonstrate strong responses to relatively mild differences in conditions. While the literature as it currently stands is suggestive of how these effects might be explained, we find nothing in the literature like the quantitative characterizations required to construct response surfaces.

The reasons for the disparity between field estimates of environmental responses on one hand and detailed studies on the other are well known in ecology. Small scale studies often produce inferences about individual effects operating under simplified conditions, particularly if these studies are carried out in the laboratory. In the wild many competing and interacting effects are present and it is often unclear how field conditions even relate to conditions constructed in the laboratory (e.g.-constant temperature treatments in a simple

landscape compared to diurnal temperature variation in the context of competition for refugia). Detailed studies typically focus on demonstrating that an effect operates, but showing that the effect generalizes to other conditions and how it is mediated by other effects is typically beyond their scope.

Detailed field studies such as mark-recapture and telemetry studies can generate estimates of many of the same effects as small-scale field and laboratory studies in the context of natural environmental variation. They can be seen as an alternative to measurements conducted in more controlled conditions, but I prefer to view them as a natural extension of the laboratory work used to first establish the existence of certain physiological or behavioral effects.

Unfortunately there are a number of barriers to using field studies broadly to quantify environmental responses. First, unlike studies that first establish the existence of effects, studies that quantify effects are often viewed as less critical. Hopefully with increasing emphasis on prediction, studies focused on quantifying effects will become more desirable. Second, even when these studies are carried out, the raw data required to compare estimates from different studies in a meta-analysis is often unavailable. Finally, for estimates of vital rates from long-term field studies or studies that cover many sites, the models required to make sense of the data are not well developed. This is the focus of my contribution.

First I demonstrated how counting process models for time-to-event data can be used to tackle the complexity of field data sets. I analyzed a simple multi-variate experimental example as well as the emigration data from the Atlantic salmon population at the West Brook. These methods are broadly applicable to the estimation of vital rates and they solve the issues of changing methods in multi-year data sets, continuous models of seasonality, and using non-constant hazard functions. More work is needed in making these methods accessible to ecologists without extensive computational experience. Work is also needed to clarify how these methods interact with the complex observational processes that are the hallmark of field experiments.

Second, I applied these methods in simplified form to a large multi-year mark-recapture data set. I used polynomial regression with interactions to construct seasonal response surfaces for survival over a range of environmental conditions. Surfaces varied strongly across seasons and were also clearly non-linear. I evaluated how environmental effects on survival correspond to the known literature on physiological and behavioral mechanisms. The simple task of comparing surfaces to the known literature makes it clear that a variety of small

targeted studies would be useful in clarifying why the environmental response surfaces look like they do. One of the strengths of this quantitative comparison is to highlight specific discrepancies between our knowledge and observations in a very directed way.

While these surfaces are valuable for understanding environmental effects on survival, they were clearly non-linear and the model was likely hiding more complex features. In general we expect environmental response surfaces to be complex when multiple unknown effects interact and are mediated by environmental conditions. To address the issue of estimating such surfaces, I suggest using local regression.

I applied one formulation of local regression, along with an evaluation of data adequacy and parameter uncertainty to the simple problem of evaluating changes in stream temperatures. The water temperature example is instructive. The simple model of change in water temperature as a function of season, forest cover, and air-water difference shows a non-linear seasonal effect, a linear air-water difference effect, and ambiguous results based on the site-level forest cover covariate (figure 28). When the interaction of forest cover with temperature difference is considered, the two are mostly independent based on the limited number of sites available (29). A site-specific model of the interaction between season and temperature difference similarly results in surfaces that could be represented by linear models with interactions. However, when aggregating across sites, the response surface for temperature difference is more complex (figure 30). The open question is whether the more complex surface aggregating across sites is simply the result of the available handful of sites or whether the pattern holds generally—or at least regionally. A general or regional pattern would be a useful component of a predictive temperature model.

In the survival model example local data are much more sparse—a single year will only have one kind of winter and the data set encompasses about twelve effective years of data. Assuming that these effects do not drift over years, we were able to estimate response surfaces for different seasons. Despite relatively broad confidence intervals these surfaces clearly depart from surfaces defined by weak priors and narrow the relevant credible intervals. In visualizing these surfaces we are able to crudely show which parts of covariate space have significant amounts of data contributing to the estimated surface or, equivalently, how confidence intervals are narrowed relative to a weak prior.

6.3   Live data, non-linearity and deep interactions

Ecological models raise a common set of challenges—varying data sources, varying process models, under-characterized non-linearity, and deep interactions, which bring with them difficulties in characterizing data adequacy and effect uncertainty in complex covariate spaces. For well studied systems it is often possible to address these challenges at the cost of creating complex custom models. I suggest a more generic approach based on tools developed in statistics to address challenges shared by many other fields.

I suggest applying a set of standard tools including time to event models and local regression for estimation and detection of non-linearity, kernel density estimation for data adequacy evaluation in complex if low-dimensional covariate spaces, and relative statistics for uncertainty evaluation. Using these relatively well-established tools is an advance over custom global models constructed from polynomial regression with interactions, as it avoids the losing proposition of trying to build the right model with what are, fundamentally, the wrong building blocks.

Even in the context of these tools there is plenty of room to specify models that address many interesting ecological questions. My goal was to solve these basic problems associated with complex field studies in enough generality to demonstrate the the tools were available to deal with complex data sources, process models, non-linearity, and deep interactions.

Appendix A
Choice of software and algorithms

The variety of mark-recapture model is well supported by software. Standard models are supported by program MARK White & Burnham (1999), which collects them in a remarkably user-friendly software package with a unified interface for model specification, comparison, and goodness-of-fit testing. At the time of this writing a variety of models have also been implement using the BUGS language in WinBUGS (Lunn *et al.*, 2000) and JAGS (Plummer, 2003). The BUGS language implementations are less user-friendly but they are amenable to iterative modification by researchers, which has allowed them to be applied to a wider class of problems (Kéry & Schaub, 2012; Royle & Dorazio, 2008). Many of these experiments might not otherwise fit into the standard mark-recapture models.

One of the BUGS-language implementations of the CJS model, which has recently become popular is the state-space formulation (Royle & Dorazio, 2008). The state-space representation tracks the state of each individual at each time point. The multi-state extension to CJS is important for our Atlantic salmon data set since individuals can depart the system by entering one of three absorbing states: death, smolting, or emigration. In addition, smolting must be preceded by emigration. In the state-space framework this type of model becomes prohibitively computationally expensive to estimate for large data sets (B. H. Letcher, Schueller, *et al.*, 2014). Instead of the state space formulation we return to the original CJS estimation strategy of summing over possible individual states to calculate the likelihood.

I do not rely on either MARK or the BUGS language implementations for the present work. Early work did rely on MARK but we were not able to obtain results for more complex models on our entire data set. Program MARK is not open source and I did not pursue a resolution to these problems. The BUGS language implementations are capable of handling larger data sets, and during the time of this project JAGS in particular improved dramatically in its robustness. We used JAGS to initially implement models for the Atlantic salmon data set, and a related project on brook trout (B. H. Letcher, Schueller, *et al.*, 2014). Our experience with the BUGS language implementation resulted in some successes but lack of convergence was a continual problem and, in more complex models, difficult to assess. I did pursue a better understanding of JAGS in particular to see if these problems were

resolvable but at the time JAGS was difficult to extend. Ultimately we were not able to *reliably* fit complex mark-recapture models for a large data set using either MARK or BUGS.

The work here depends on the Stan (Stan Development Team, 2014) library and interfaces. Stan supports a language superficially similar to BUGS but its architecture is by design more scalable. Briefly, rather than creating a DAG which is then used to conditionally sample from the posterior, Stan's specification is parsed into a C++ program which defines the C++ density calculation. The C++ program is then compiled and a variety of algorithms can be used to optimize or sample from the posterior. This work relies on the Hamiltonian Monte Carlo algorithm in Stan for full Bayesian inference, which has excellent properties for sampling from a posterior with high and varying correlation among parameters.

With new software there is always the concern that some results may be wrong simply due to the software. To partially address that concern, I estimated some models that did not include a random effect structure with optimization (BFGS-L) in Stan, as well as sampling. This tests that spurious results are generally not introduced by the two more complex components of the Stan library which sampling relies on—automatic gradient calculation and the HMC/NUTS algorithm. The results were generally very close and I do not discuss these further here.

Stan was somewhat slow in some of our very complex models that included three-dimensional spline structures with hundreds of knot points. The run-times for these models approached two or three days, which made model building and evaluation time-consuming. The relevant calculations benefit from sparse matrix algebra which was lacking in Stan. To solve this problem I programmed a small addition to Stan in C++ which made it possible to use a sparse model matrix to represent the multi-dimensional spline and cut computation down to hours. The changes have since been incorporated into Stan. While the C++ API for Stan is not yet well documented I expect it to improve quickly as the team is very committed to an open development model.

Appendix B

Commented example codes

The code presented here is available from the author as well as online (Sakrejda, n.d.) and retrieving up-to-date sources is preferable to working with this code. This code is provided as a reference for the computational issues involved in coding the statistical models described in the main text.

Comments on modified CJS code

The modified CJS model used in the Atlantic salmon analysis is relatively standard as both the survival and recapture calculations used a pre-computed model matrix. The general model appears here with comments, the complete code is available online.

The smolt model was very limited and specific. I abstracted the calculation of the smolt probability to this function. Smolt probability is zero in all non-spring seasons and below age 2^+ whereas α is indexed to $a[1]$ for age 1^+ and to $a[2]$ for all older ages.

```
functions {
real beta_smolt(vector alpha, int a, int s) {
  if (s!=1) return 0.0;
  if (a<=1) return 0.0;
  if ((a-1) < num_elements(alpha))
    return inv_logit(alpha[a-1]);
  else
    return inv_logit(alpha[num_elements(alpha)]);
}
}
```

The indexing required for the long form of mark-recapture data here is relatively involved but it saves on empty matrix entries relative to the wide form and allows the data to be kept in a format relatively convenient for visual inspection. Ambiguous individuals are those without a known fate and ambiguous rows are rows of data where the fate of the individual is not known. These are typically all rows after the last capture if there is no

subsequent antenna-detected emigration or smolt trap recapture. Unambiguous fates are more convenient as those individuals can be censored at their final observation and their likelihood calculation is simpler. Season and smolt-season are different variables as some smolts were caught in the smolt trap after the official end of spring, but they were only a handful and so were lumped with the spring smolts.

The somewhat generic priors we apply require us to separate survival and recapture events into fixed and random effects (Cauchy for recapture and normal for survival). The data relating survival and recapture parameters to observations (and environmental conditions) is passed in as a sparse-format model matrix (CSR) which does not have a natural data type in the Stan language (yet). Instead of using a single type, each model matrix is represented by the number of nonzero entries (NNZ), the number of rows (m), the number of columns (n), the values of non-zero entries (w), the row-index of non-zero entries (v), and two additional vectors (u, z). These are documented further in the section describing sparse matrix multiplication implemented in Stan.

```
data {
  int n_rows;
  int n_censored;
  int n_ambiguous_individuals;

  int n_unambiguous_rows;
  int n_ambiguous_rows;
  int n_individuals;
  int ambiguous_rows[n_ambiguous_individuals,2];
  int unambiguous_rows[n_unambiguous_rows];

  int season[n_rows];
  int smolt_season[n_rows];
  int age_year[n_rows];
  vector<lower=0, upper=10>[n_rows] scaled_interval;
  int cjs_classification[n_rows];

  int n_p_effects;
```

```
    int n_p_cauchy_effects;
    int n_p_fixed_effects;
    int n_phi_effects;
    int n_phi_random_effects;
    int n_phi_fixed_effects;

//  matrix[n_rows,n_p_effects] p_X;
    int p_X_NNZ;
    int p_X_m;
    int p_X_n;
    vector[p_X_NNZ] p_X_w;
    int p_X_v[p_X_NNZ];
    int p_X_u[p_X_m+1];
    int p_X_z[p_X_m];

//  matrix[n_rows,n_phi_effects] phi_X;
    int phi_X_NNZ;
    int phi_X_m;
    int phi_X_n;
    vector[phi_X_NNZ] phi_X_w;
    int phi_X_v[phi_X_NNZ];
    int phi_X_u[phi_X_m+1];
    int phi_X_z[phi_X_m];
}
```

In the transformed data section Stan makes it convenient to carry out some deterministic calculations on data only—before running the estimation algorithm. First, we need to know, based on the data, how long of a vector we will need to keep for calculations on individuals with non-censored histories. This is determined from the number of non-capture rows after the last capture for each individual. We start with a zero length buffer and check whether it is long enough for each individual. When it is (inevitably) found to be too short, we update the value to be long enough for the offending individual. After the loop covers all

individuals we store the length of the buffer so we can pre-allocate it and avoid creating a new one for each individual. Second, we check whether the number of fixed effects in the recapture model matches the number of random effects and total number of effects. These values could all be calculated from two out of three but instead we use it as a check on the R-level calculations.

```
transformed data {
  int ambig_row_calc_buffer;
  int m;

  ambig_row_calc_buffer <- 0;
  for ( i in 1:n_ambiguous_individuals ) {
    m <- ambiguous_rows[i,2] - ambiguous_rows[i,1] + 1;
    if ( m > ambig_row_calc_buffer ) {
      ambig_row_calc_buffer <- m;
    }
  }


  if (n_p_effects != n_p_cauchy_effects + n_p_fixed_effects)
    print("n_p_effects != n_p_cauchy_effects + n_p_fixed_effects");


}
```

In the following section we declare the survival, recapture, smolt and emigration parameters. There are two smolt parameters (one for 1^+ and one for all older fish), four emigration parameters (one per season, based on the exploratory data analysis of emigration). The survival parameters are divided into random and fixed effects so that they can be assigned priors in groups below, with one variance parameter for the random effects. The limits on the survival parameters limit the model to the portion of parameter space where the logistic transformation is still sensitive to parameter values. Parameters much outside of $[-7, 7]$ are essentially zero or one after the inverse logit transform, so we limit the values they can take to simplify analysis. The recapture parameters follow the same pattern. The recapture parameters have different bounds as they are later transformed to induce a Cauchy distribution on the random effects. This change is important for the Cauchy random effects as

sampling from a Cauchy distribution directly with an MCMC algorithm is in general not effective.

```
parameters {
  vector<lower=-7, upper=7>[n_phi_random_effects] phi_random_betas_raw;
  vector<lower=-7, upper=7>[n_phi_fixed_effects] phi_betas_raw;
  real<lower=0, upper=10>  phi_sd;

  vector<lower=-pi()/2, upper=pi()/2>[n_p_cauchy_effects] p_cauchy_betas_raw;
  vector<lower=-7, upper=7>[n_p_effects-n_p_cauchy_effects]  p_fixed_betas;
  real<lower=0, upper=3> beta_p_occasion_sd;

  vector<lower=-15, upper=5>[2] beta_alpha_seasonal; // smolt
  vector<lower=-15, upper=5>[4] beta_rho_seasonal;   // emigrate
}
```

In the transformed parameters the vectors of fixed-effect and random-effect betas are combined to create a single vector for survival and a single vector for recapture. The recapture random effects are transformed such that their induced distribution is Cauchy with an estimated standard deviation.

```
transformed parameters {
  vector[n_phi_effects] phi_betas;
  vector[n_p_effects] p_betas;

  for ( i in 1:n_phi_fixed_effects) {
    phi_betas[i] <- phi_betas_raw[i];
  }
  for ( i in (n_phi_fixed_effects+1):n_phi_effects ) {
    phi_betas[i] <- phi_sd * phi_random_betas_raw[i-n_phi_fixed_effects];
  }

  for ( i in 1:n_p_fixed_effects) {
    p_betas[i] <- p_fixed_betas[i];
```

```
  }
  for ( i in (n_p_fixed_effects+1):n_p_effects ) {
    p_betas[i] <- beta_p_occasion_sd *
    tan(p_cauchy_betas_raw[i-n_p_fixed_effects]);   // centered cauchy random effects.
  }


}
```

The model block is complex so we break it down into sections below. The first portion is declaring variables and types. We declare buffers (A, A^*, B, and B^*) and their sizes based on the calculation from the transformed data block. We also declare a per-row likelihood value. The priors for recapture are $N(0, 1)$, which places the main mass of the distribution between a survival probability of 0.1 and 0.9 for a single transformed coefficient and wider for the result of a linear model. The specific formulation below has no random effects and therefore the same priors for all types of survival effects. The recapture effects also have the weak $N(0, 1)$ prior for fixed effects and an implied uniform prior for the random effects which are then transformed to induce a Cauchy distribution.

```
model {
  real row_lp[n_rows];
  int k;
  int start_row;
  int stop_row;
  real A[ambig_row_calc_buffer];
  real A_star[ambig_row_calc_buffer];
  real B[ambig_row_calc_buffer];
  real B_star[ambig_row_calc_buffer];

  vector[n_rows] phi;
  vector[n_rows] p;
  vector[n_rows] alpha;
  vector[n_rows] rho;
```

```
// survival priors:
phi_betas_raw ~ normal(0,1);
phi_random_betas_raw ~ normal(0,1);

// most of the mass below sd=1, still allows extreme events
phi_sd ~ gamma(2,10);

// recapture priors:
// p_cauchy_betas_raw  ~unif()  // implied!
p_fixed_betas ~ normal(0,1);
beta_p_occasion_sd ~ cauchy(0,.5);

// smolting priors:
beta_alpha_seasonal ~ cauchy(-1,3);

// recapture priors:
beta_rho_seasonal ~ normal(-5,3);
```

Continuing the model block, the survival and recapture model matrices (in sparse matrix format) are multiplied by their respective parameter to produce per-observation link-scale values which are then transformed (and/or scaled by per-observation interval duration) to produce per-observation survival/recapture/smolt/emigrate event probabilities.

```
// construct per-event probs:
phi <- sparse_multiply_csr(phi_X_m, phi_X_n, phi_X_w,
        phi_X_v, phi_X_u, phi_X_z, phi_betas);
p <- sparse_multiply_csr(p_X_m, p_X_n, p_X_w, p_X_v, p_X_u, p_X_z, p_betas);

for ( i in 1:n_rows ) {
  phi[i] <- pow(inv_logit(phi[i]), scaled_interval[i]);
  p[i] <- inv_logit(p[i]);
  alpha[i] <- beta_smolt(beta_alpha_seasonal, age_year[i], smolt_season[i]);
  rho[i] <- inv_logit(beta_rho_seasonal[season[i]]);
}
```

Finally, the event probabilities are combined based on the actual observed event type to produce a per-observation likelihood. For events which should *not* be included in the model there are guard statements included in this calculation. For all observations between the first and last the calculation is straightforward based on the known combinations of events leading to an observation. The calculation for individuals with uncensored histories ("ambiguous") have a second recursive calculation to compute the probability of each non-observation based on the parameters and the history of observation attempts. This recursive calculation re-uses the single buffer instantiated at the top of the model block.

```
// Likelihood calculation:
for ( i in 1:n_unambiguous_rows ) {
  k <- unambiguous_rows[i];
  if (cjs_classification[k] == 0) {
    print("STOP, row ", k, " should not be evaluated.");
    row_lp[k] <- 0.0;
  }
  if (cjs_classification[k] == 1) {
    row_lp[k] <- log(phi[k]*(1-rho[k])*(1-alpha[k])*p[k]);
  }
  if (cjs_classification[k] == 2) {
    row_lp[k] <- log(phi[k]*(1-rho[k])*(1-alpha[k])*(1-p[k]));
  }
  if (cjs_classification[k] == 3) {
    row_lp[k] <- log(phi[k]*rho[k]);
  }
  if (cjs_classification[k] == 4) {
    row_lp[k] <- log(phi[k]*rho[k]*alpha[k]);
  }
  if (cjs_classification[k] == 5) {
    row_lp[k] <- log(phi[k]*(1-rho[k])*(1-alpha[k]));
  }
  if (cjs_classification[k] == 6) {
    print("STOP, row ", k, " should not be evaluated here.");
```

136

```
  }
  if (cjs_classification[k] == 7) {
    print("STOP, row ", k, " should not be evaluated.");
    row_lp[k] <- 0.0;
  }
  increment_log_prob(row_lp[k]);
}

for ( i in 1:(n_ambiguous_individuals)) {
  for ( j in 1:ambig_row_calc_buffer ) {
    A[j] <- 0.0;
    B[j] <- 0.0;
    A_star[j] <- 0.0;
    B_star[j] <- 0.0;
  }
  start_row <- ambiguous_rows[i,1];
  stop_row <- ambiguous_rows[i,2];
  for ( j in start_row:stop_row ) {
    k <- j - start_row + 1;
    if (cjs_classification[j] != 6) {
      print("STOP, row ", j, " should not be evaluated here.");
    }
    if (j == start_row) {
      A[k] <- 1.0;
      A_star[k] <- 1.0;
    } else {
      A[k] <- phi[j-1]*(1-p[j-1]);
    }
    B[k] <- (1-phi[j]);

    A_star[k] <- A_star[k] * A[k];
    B_star[k] <- A_star[k] * B[k];
```

```
    }
    increment_log_prob(log(sum(B_star)));
  }
```

Comments on continuous-time CJS code

```
functions {
  real protected_normal_cdf(real x, real mu, real sd) {
    if ((x-mu)/sd > 10.0)
      return(1.0);
    else if ((x-mu)/sd < 10.0)
      return(0.0);
    else
      return normal_cdf(x,mu,sd);
  }

  real lambda_radial(vector theta, real t, vector process_knot_centers,
    real knot_scale
  ) {
    real lambda;
    int K;
    K <- num_elements(theta);
    lambda <- 0.0;
    for(k in 1:K) {
      lambda <- lambda +  theta[k] *
        exp(normal_log(t, process_knot_centers[k], knot_scale));
    }
    return lambda;
  }

  real Lambda_radial(vector theta, real s1, real t,
```

```
vector process_knot_centers, real knot_scale) {
  real lambda;
  int K;
  K <- num_elements(theta);
  lambda <- 0.0;
  for(k in 1:K) {
    lambda <- lambda +  theta[k] * (
      normal_cdf(t, process_knot_centers[k], knot_scale) -
      normal_cdf(s1, process_knot_centers[k], knot_scale)
    );
  }
  return lambda;
}

real time_to_event(
  int event_class, real s1, real s2,
  vector theta_D, vector theta_E,
  vector process_knot_centers, real knot_scale
) {
  real log_prob;
  real lambda_D;
  real lambda_E;
  real Lambda_D;
  real Lambda_E;
  if (event_class == 0) { // non-event, no effect
    return 0.0;
  }
  if (event_class == 1) { // no emigration, no death
    Lambda_D <- Lambda_radial(
    theta_D, s1, s2, process_knot_centers, knot_scale);
    Lambda_E <- Lambda_radial(
    theta_E, s1, s2, process_knot_centers, knot_scale);
    log_prob <- -(Lambda_D+Lambda_E);
```

```
    return log_prob;
  }
  if (event_class == 2) { // known emigration, no death
    Lambda_D <- Lambda_radial(
    theta_D, s1, s2, process_knot_centers, knot_scale);
    lambda_E <- lambda_radial(
    theta_E,     s2, process_knot_centers, knot_scale);
    Lambda_E <- Lambda_radial(
    theta_E, s1, s2, process_knot_centers, knot_scale);
    log_prob <- log(lambda_E) - (Lambda_D+Lambda_E);
    return log_prob;
  }
  if (event_class == 3) { // no emigration, known death
    lambda_D <- lambda_radial(
    theta_D,     s2, process_knot_centers, knot_scale);
    Lambda_D <- Lambda_radial(
    theta_D, s1, s2, process_knot_centers, knot_scale);
    Lambda_E <- Lambda_radial(
    theta_E, s1, s2, process_knot_centers, knot_scale);
    log_prob <- log(lambda_D) - (Lambda_D+Lambda_E);
    return log_prob;
  }
  if (event_class == 4) { // no emigration, death on interval
    print("Stop, this likelihood can not be calculated.")
    log_prob <- not_a_number();
    return log_prob;
  }
  if (event_class == 5) { // emigration on interval, no death
    print("Stop, this likelihood can not be calculated.")
    log_prob <- not_a_number();
    return log_prob;
  }
  if (event_class == 6) { // one of emigration or death
```

```
    lambda_D <- lambda_radial(
    theta_D,    s2, process_knot_centers, knot_scale);
    Lambda_D <- Lambda_radial(
    theta_D, s1, s2, process_knot_centers, knot_scale);
    lambda_E <- lambda_radial(
    theta_E,    s2, process_knot_centers, knot_scale);
    Lambda_E <- Lambda_radial(
    theta_E, s1, s2, process_knot_centers, knot_scale);
    log_prob <- log(lambda_D+lambda_E) - (Lambda_D+Lambda_E);
    return log_prob;
  }
  return not_a_number();
}
real smooth_recapture_event(real s1, real s2,
vector p, vector occ_centers, vector occ_spreads) {
  real log_prob;
  int n_occ;
  n_occ <- num_elements(p);
  log_prob <- 0.0;
  for (i in 1:n_occ) {
    log_prob <- log_prob + log(
      1 - p[i] * (protected_normal_cdf(s2, occ_centers[i], occ_spreads[i]) -
        protected_normal_cdf(s1, occ_centers[i], occ_spreads[i]) )
    );
  }
  return log_prob;
}

}

data {
  int n_rows; // number of rows/observations/events
  int n_individuals; // number of individuals/units
```

```
int n_known_s2;
int idx_known_s2[n_known_s2];
int n_unknown_s2;
int idx_unknown_s2[n_unknown_s2];

// class is defined by code in function above.
int idx_event_class[n_rows];
vector[n_rows] s1;
// Known or best guess values for interval end.
vector[n_rows] s2;

// Both death and emigration share the same number of knots and the
int n_process_knots;
// same knot centers/spreads to avoid mobile effects.
vector[n_process_knots] process_knot_centers;

int n_occasions;
vector[n_occasions] start_occ;
vector[n_occasions] stop_occ;


// class is defined by code in function above.
int idx_recapture_class[n_rows];
int idx_occ[n_rows];

// for failed recaptures with unknown state, this data is used
// as the "s1" for the recapture event calculation.  It is an
// artifical "s1" which is really the day following the last day
// of the previous recapture occasion.
vector[n_rows] recapture_s1;
int n_p_effects;
int n_p_cauchy_effects;
```

```
  int n_p_fixed_effects;
  matrix[n_rows,n_p_effects] p_X;
  matrix[n_occasions, n_p_effects] p_X_occ;

  int start_day;  // For generated quantities, lambdas are calculated
  int stop_day;    // for each day of interest.  These are the start/stop
                   // days (with zero absolute reference, like the s1/s2/t
                   // variables.
}
transformed data {
  real knot_scale;
  real knot_weight_scale;
  vector[n_process_knots-1] interknot;
  vector[n_occasions] occ_centers;
  vector[n_occasions] occ_spreads;

  for(k in 2:n_process_knots)
    interknot[k-1] <- process_knot_centers[k] - process_knot_centers[k-1];
  knot_scale <- mean(interknot)/2.0;  // Based on interknot distance,
                                      // this scale makes for a flat
                                      // function when weights are
                                      // equal.
  knot_weight_scale <- 1/exp(normal_log(0,0,knot_scale));
    // Scaling thetas by this makes the function about the height of
    // thetas near knot centers.

  for(i in 1:n_occasions) {
    occ_centers[i] <- (start_occ[i]+stop_occ[i])/2;
    occ_spreads[i] <- (stop_occ[i]-start_occ[i])/4;
    if (occ_spreads[i] < 1)
      occ_spreads[i] <- 1;
  }
}
```

```
parameters {
  // All parameters on YEARLY scale.
  vector<lower=0>[n_unknown_s2] t2;
  vector<lower=0>[n_process_knots] theta_D;
  vector<lower=0>[n_process_knots] theta_E;

  // Recapture, per-occasion and spline parameters:
  vector<lower=-pi()/2, upper=pi()/2>[n_p_cauchy_effects] p_cauchy_betas_raw;
  vector<lower=-7, upper=7>[n_p_effects-n_p_cauchy_effects] p_fixed_betas;
  real<lower=0, upper=3> p_occasion_sd;

}

transformed parameters {
 vector[n_rows] t;  // parameters for unkonwn event times;

 vector[n_p_effects] p_betas;
 vector[n_occasions] p_occ; // per-occsaion recapture probability

  for ( i in 1:n_known_s2 ) {
    t[idx_known_s2[i]] <- s2[idx_known_s2[i]];
  }
  for ( i in 1:n_unknown_s2) {
   // t2 survival in years TO days.
    t[idx_unknown_s2[i]] <- s1[idx_unknown_s2[i]] + t2[i]*365;
  }

  for ( i in 1:n_p_fixed_effects) {
    p_betas[i] <- p_fixed_betas[i];
  }
  for ( i in (n_p_fixed_effects+1):n_p_effects ) {
    // centered cauchy random effects.
```

```stan
    p_betas[i] <- p_occasion_sd * tan(p_cauchy_betas_raw[i-n_p_fixed_effects]);
  }

  // Per-occasion recapture probabilities:
  p_occ <- p_X_occ * p_betas;
  for(i in 1:n_occasions) {
    p_occ[i] <- inv_logit(p_occ[i]);
  }

}


model {
  vector[n_rows] lambda_D; // per-observation intensity function value.
  vector[n_rows] Lambda_D; // per-observation mean function value.
  vector[n_rows] lambda_E; // per-observation intensity function value.
  vector[n_rows] Lambda_E; // per-observation mean function value.

  vector[n_rows] p; // per-row recapture probability

  vector[n_rows] log_lik;

  t2 ~ gamma(2,1);                // YEARLY BASIS STUFF!!!!!
  theta_D ~ gamma(2,1*knot_weight_scale);
  theta_E ~ gamma(2,1*knot_weight_scale);

  p_fixed_betas ~ normal(0,1);
  p_occasion_sd ~ gamma(2,1);

  //p <- sparse_multiply_csr(p_X_m, p_X_n, p_X_w, p_X_v, p_X_u, p_X_z, p_betas);
  p <- p_X * p_betas;

  for(i in 1:n_rows) {
    p[i] <- inv_logit(p[i]);
```

```
  }

  for ( i in 1:n_rows ) {
    log_lik[i] <- time_to_event(
      idx_event_class[i], s1[i], t[i], theta_D/365*knot_weight_scale,
      theta_E/365*knot_weight_scale,
      process_knot_centers, knot_scale);
//  if (idx_recapture_class[i]==0)
//    log_lik[i] <- log_lik[i] + 0.0;
    if (idx_recapture_class[i]==1)
      log_lik[i] <- log_lik[i] + log(p[i]);
    if (idx_recapture_class[i]==2)
      log_lik[i] <- log_lik[i] + log(1-p[i]);
    if (idx_recapture_class[i]==3)
      log_lik[i] <- log_lik[i] +
        smooth_recapture_event(recapture_s1[i], t[i],
          p_occ, occ_centers, occ_spreads);
    increment_log_prob(log_lik[i]);
  }

}


generated quantities {
  vector[n_rows] event_log_lik;
  vector[n_rows] recap_log_lik;
  vector[n_rows] p;
  p <- p_X * p_betas;

  for(i in 1:n_rows) {
    p[i] <- inv_logit(p[i]);
  }
  for ( i in 1:n_rows ) {
```

```
    event_log_lik[i] <- time_to_event(
      idx_event_class[i], s1[i], s2[i], theta_D/365*knot_weight_scale,
      theta_E/365*knot_weight_scale,
      process_knot_centers, knot_scale);
    if (idx_recapture_class[i]==1)
      recap_log_lik[i] <- log(p[i]);
    if (idx_recapture_class[i]==2)
      recap_log_lik[i] <- log(1-p[i]);
    if (idx_recapture_class[i]==3)
      recap_log_lik[i] <- smooth_recapture_event(recapture_s1[i], s2[i],
      p_occ, occ_centers, occ_spreads);
  }


}
```

Comments on multi-state homogeneous Poisson process code

Unlike the two versions of the CJS model, this code is specific to the conceptual model and
is an approximation. While it does consider forward and back transitions in the three-state
system, it does not consider multi-step transitions. In some systems multi-step transitions
may be important. The single-transition function below is complicated by a branch point
because the integral can not be calculated with the same method at all parameter values.
When the forward and reverse transitions have equal rates the integral simplifies to a different
form but the original formula is undefined.

```
functions {
  real no_transition(real[] s, vector lambda, int a_idx, int b_idx) {
    real log_prob;
    log_prob <- -((lambda[a_idx] + lambda[b_idx])*(s[2]-s[1]));
    return(log_prob);
```

```
  }

  real one_transition(real[] s, vector lambda,
  int e_idx, int c_idx, int r1_idx, int r2_idx) {
    real log_prob;
    real log_a;
    real b;
    real d;
    log_a <- log(lambda[e_idx]) +
         (lambda[e_idx]+lambda[c_idx])*s[1] -
         (lambda[r1_idx]+lambda[r2_idx])*s[2];
    b <- lambda[e_idx] + lambda[c_idx] -
 lambda[r1_idx] - lambda[r2_idx];
    if (b == 0) {
      log_prob <- log_a + log(s[2]-s[1]);
      return(log_prob);
    } else {
      d <- exp(-b*s[1]) - exp(-b*s[2]);
      log_prob <- log_a + log(d/b);
      return(log_prob);
    }
  }
}
```

Simplifying the approximate model to only consider the no-transition and single-transition cases allows us to simplify the data by keeping track of the index of the forward transition (e), the complementary transition to the forward transition (c), which did not occur, and the two reverse transitions (r1, r2), which also did not occur.

```
data {
  // structure:
  int n_obs;
```

```
      int n_transitions;
      int n_groups;

      // times:
      real<lower=0> s[n_obs,2];

      // groups:
      int<lower=1, upper=2>  transition_type[n_obs]; // self=1, other=2
      int<lower=1, upper=n_groups>  group[n_obs]; // 1:n_groups;

      // Index lambdas: event, competing, reverse 1, reverse 2;
      int<lower=1, upper=n_transitions> e_idx[n_obs];
      int<lower=1, upper=n_transitions> c_idx[n_obs];
      int<lower=0, upper=n_transitions> r1_idx[n_obs];
      int<lower=0, upper=n_transitions> r2_idx[n_obs];


}
```

The parameters here are simple with one set of lambdas per treatment group, and one lambda within group for each of the available transitions (e, c, r1, and r2). The model block is also simple, adding one of two possible increments—no transition observed, or transition observed—to the log probability for each data point.

```
parameters {
  matrix<lower=0>[n_transitions,n_groups] lambda;
}

model {
  real log_prob;

  for (i in 1:n_groups) {
```

```
    col(lambda, i) ~ gamma(2,1);
  }

  for(i in 1:n_obs) {
    if (transition_type[i]==1) {
      log_prob <- no_transition(s[i],
      col(lambda, group[i]), e_idx[i], c_idx[i]);
    } else {
      log_prob <- one_transition(s[i],
      col(lambda, group[i]), e_idx[i], c_idx[i], r1_idx[i], r2_idx[i]);
    }
    increment_log_prob(log_prob);
  }

}
```

REFERENCES

Anderson, T. C. & MacDonald, B. P. (1978). *A Portable Weir for Counting Migrating Fishes in Rivers.* Technical Report 733. Canada Fisheries and Marine Service.

Armstrong, J. D., Braithwaite, V. A. & Fox, M. (1998). The Response of Wild Atlantic Salmon Parr to Acute Reductions in Water Flow. *Journal of Animal Ecology*, 67, 292–297.

Arnold, G. P., Webb, P. W. & Holford, B. H. (1991). The Role of the Pectoral Fins in Station-Holding of Atlantic Salmon Parr (Salmo Salar L.) *Journal of Experimental Biology*, 156, 625–629.

Austin, M. P. (2002). Spatial prediction of species distribution: an interface between ecological theory and statistical modelling. *Ecological modelling*, 157, 101–118.

Bacon, P. J., Gurney, W. S. C., Jones, W., McLaren, I. S. & Youngson, A. F. (2005). Seasonal Growth Patterns of Wild Juvenile Fish: Partitioning Variation Among Explanatory Variables, Based on Individual Growth Trajectories of Atlantic salmon (*Salmo salar*) Parr. *Journal of animal ecology*, 74, 1–11.

Bal, G., Rivot, E., Prévost, E., Piou, C. & Balinière, J. L. (2011). Effect of Water Temperature and Density of Juvenile Salmonids on Growth of Young-of-the-Year Atlantic Salmon *Salmo salar*. *Journal of Fish Biology*,

Barbet-Massin, M., Thuiller, W. & Jiguet, F. (2010). How Much Do we Overestimate Future Local Extinction Rates when Restricting the Range of Occurrence Data in Climate Suitability Models? *Ecography*, 33, 878–886.

Bardonnet, A. & Baglinière, J.-L. (2000). Freshwater Habitat of Atlantic Salmon (*Salmo salar*). *Canadian Journal of Fisheries and Aquatic Sciences*, 57, 497–506.

Barker, R. J. (1997). Joint Modeling of Live-Recapture, Tag-Resight, and Tag-Recovery Data. *Biometrics*, 53, 666–677.

Berg, O. K. & Bremset, G. (1998). Seasonal Changes in the Body Composition of Young Riverine Atlantic Salmon and Brown Trout. *Journal of Fish Biology*, 52, 1272–1288.

Biro, P. A., Post, J. R. & Parkinson, E. A. (2003). From Individuals to Populations: Prey Fish Risk-Taking Mediates Mortality in Whole-System Experiments. *Ecology*, 84, 2419–2431.

151

Breau, C., Cunjak, R. A. & Bremset, G. (2007). Age-Specific Aggregation of Wild Juvenile Atlantic Salmon (*Salmo salar*) at Cool Water Sources During High Temperature Events. *Journal of Fish Biology*, 71, 1179–1191.

Brodeur, R. D. & Wilson, M. T. (1996). A Review of the Distribution, Ecology and Population Dynamics of Age-0 Walleye Pollock in the Gulf of Alaska. *Fisheries Oceanography*, 5, 148–166.

Castro-Santos, T., Haro, A. & Walk, S. (1996). A Passive Integrated Transponder (PIT) Tag System for Monitoring Fishways. *Fisheries Research*, 28, 253–261.

Caswell, H. (2001). *Matrix Population Models*. Sinauer Associates.

Chandler, R. B. & Clark, J. D. (2014). Spatially Explicit Integrated Population Models. *Methods in Ecology and Evolution*, 5, 1351–1360.

Chen, I.-C., Hill, J. K., Ohlemüller, R., Roy, D. B. & Thomas, C. D. (2011). Rapid Range Shifts of Species Associated with High Levels of Climate Warming. *Science*, 333, 1024–1026.

Clark, J. S., Carpenter, S. R., Barber, M., Collins, S., Dobson, A., Foley, J. A., Lodge, D. M., Pascual, M., Pielke Jr, R. & Pizer, W. (2001). Ecological Forecasts: an Emerging Imperative. *Science*, 293, 657–660.

Clotfelter, E. D., Curren, L. J. & Murphy, C. E. (2006). Mate Choice and Spawning Success in the Fighting Fish *Betta splendens*: the Importance of Body Size, Display Behavior and Nest Size. *Ethology*, 112, 1170–1178.

Conradt, L. & Roper, T. J. (2000). Activity Synchrony and Social Cohesion: a Fission-Fusion Model. *Proceedings of the Royal Society of London B: Biological Sciences*, 267, 2213–2218.

Contamin, R. & Ellison, A. M. (2009). Indicators of Regime Shifts in Ecological Systems: What Do we Need to Know and When Do we Need to Know it. *Ecological Applications*, 19, 799–816.

Cormack, R. M. (1964). Estimates of Survival from the Sighting of Marked Animals. *Biometrika*, 51, 429–438.

Cox, D. R. & Oakes, D. (1984). *Analysis of Survival Data*. Vol. 21. Monographs on Statistics and Applied Probability. CRC Press.

Crone, E. E., Ellis, M. M., Morris, W. F., Stanley, A., Bell, T., Bierzychudek, P., Ehrlén, J., Kaye, T. N., Knight, T. M., Lesica, P., Oostermeijer, G., Quintana-Ascencio, P. F., Ticktin, T., Valverde, T., Williams, J. L., Doak, D. F., Ganesan, R., McEachern, K., Thorpe, A. S. & Menges, E. S. (2013). Ability of Matrix Models to Explain the Past and Predict the Future of Plant Populations. *Conservation Biology*, 27, 968–978.

Cunjak, R. A. (1988). Behaviour and Microhabitat of Young Atlantic Salmon (*Salmo salar*) During Winter. *Canadian Journal of Fisheries and Aquatic Sciences*, 45, 2156–2160.

Cunjak, R. A., Prowse, T. D. & Parrish, D. L. (1998). Atlantic Salmon (*Salmo salar*) in Winter: The Season of Parr Discontent? *Canadian Journal of Fisheries and Aquatic Sciences*, 55, 161–180.

D'Andrea, W. J., Huang, Y., Fritz, S. C. & Anderson, N. J. (2011). Abrupt Holocene Climate Change as an Important Factor for Human Migration in West Greenland. *Proceedings of the National Academy of Sciences*, 108, 9765–9769.

Darroch, J. N. (1958). The Multiple-Recapture Census: I Estimation of a Closed Population. *Biometrika*, 45, 343–359.

Darroch, J. N. (1959). The Multiple-Recapture Census: II Estimation When There is Immigration or Death. *Biometrika*, 46, 336–351.

Davidson, R. S., Letcher, B. H. & Nislow, K. H. (2010). Drivers of Growth Variation in Juvenile Atlantic Salmon (*Salmo salar*): An Elasticity Analysis Approach. *Journal of Animal Ecology*, 79, 1113–1121.

Dionne, M. & Dodson, J. J. (2002). Impact of Exposure to a Simulated Predator (*Mergus merganser*) on the Activity of Juvenile Atlantic Salmon (*Salmo salar*) in a Natural Environment. *Canadian journal of zoology*, 80, 2006–2013.

Downie, A.-L., Numers, M. von & Boström, C. (2013). Influence of Model Selection on the Predicted Distribution of the Seagrass, *Zostera marina*. *Estuarine, Coastal and Shelf Science*, 121-122, 8–19.

Dullinger, S., Gattringer, A., Thuiller, W., Moser, D., Zimmermann, N. E., Guisan, A., Willner, W., Plutzar, C., Leitner, M., Mang, T., Caccianiga, M., Dirnbock, T., Ertl, S., Fischer, A., Lenoir, J., Svenning, J.-C., Psomas, A., Schmatz, D. R., Silc, U., Vittoz, P. & Hulber, K. (2012). Extinction Debt of High-Mountain Plants Under Twenty-First-Century Climate Change. *Nature Climate Change*, 2, 619–622.

Eilers, P. H. & Marx, B. D. (2010). Splines, Knots, and Penalties. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2, 637–653.

Elliott, J. M. (1991). Tolerance and Resistance to Thermal Stress in Juvenile Atlantic Salmon, *Salmo salar*. *Freshwater Biology*, 25, 61–70.

Elliott, J. M. & Elliott, J. A. (2010). Temperature Requirements of Atlantic Salmon *Salmo salar*, Brown Trout *Salmo trutta* and Arctic Charr *Salvelinus alpinus*: Predicting the Effects of Climate Change. *Journal of Fish Biology*, 77, 1793–1817.

Elliott, J. M. & Hurley, M. A. (1997). A Functional Model for Maximum Growth of Atlantic Salmon Parr, *Salmo salar*, from Two Populations in Northwest England. *Functional Ecology*, 11, 592–603.

Ellner, S. P. & Rees, M. (2006). Integral projection models for species with complex demography. *The American Naturalist*, 107, 410–428.

Enders, E. C., Buffin-Belanger, T., Boisclair, D. & Roy, A. G. (2005). The Feeding Behaviour of Juvenile Atlantic Salmon in Relation to Turbulent Flow. *Journal of Fish Biology*, 66, 242–253.

Fieberg, J. & DelGiudice, G. D. (2009). What Time is it? Choice of Time Origin and Scale in Extended Proportional Hazards Models. *Ecology*, 90, 1687–1697.

Finstad, A. G., Berg, O. K., Forseth, T., Ugedal, O. & Naesje, T. F. (2010). Adaptive Winter Survival Strategies: Defended Energy Levels in Juvenile Atlantic Salmon Along a Latiduinal Gradient. *Proceedings of the Royal Society B: Biological Sciences*, 277, 1113–1120.

Finstad, A. G., Ugedal, O., Forseth, T. & Naesje, T. F. (2004). Energy-Related Juvenile Winter Mortality in a Northern Population of Atlantic Salmon (*Salmo salar*). *Canadian Journal of Fisheries and Aquatic Sciences*, 61, 2358–2368.

Forseth, T., Hurley, M. A., Jensen, A. J. & Elliott, J. M. (2001). Functional Models for Growth and Food Consumption of Atlantic Salmon Parr, *Salmo salar*, from a Norwegian River. *Freshwater Biology*, 46, 173–186.

Freedman, D. A. (2008). Survival Analysis: A Primer. *The American Statistician*, 62, 110–119.

Fuller, A. & Earley, R. (2015). Interspecific Assessment of Sexual Signals: Can Heterospecifics Choose the Best Hosts?

Gardiner, W. R. & Geddes, P. (1980). The Influence of Body Composition on the Survival of Juvenile Salmon. *Hydrobiologia*, 69, 67–72.

Gries, G. & Letcher, B. H. (2002). A Night-Seining Technique for Sampling Juvenile Atlantic Salmon in Streams. *American Journal of Fisheries Management*, 22, 595–601.

Gries, G. & Juanes, F. (1998). Microhabitat Use by Juvenile Atlantic Salmon (*Salmo salar*) Sheltering During the Day in Summer. *Canadian Journal of Zoology*, 76, 1441–1449.

Guisan, A. & Thuiller, W. (2005). Predicting Species Distribution: Offering More Than Simple Habitat Models. *Ecology Letters*, 8, 993–1009.

Guisan, A. & Zimmermann, N. E. (2000). Predictive Habitat Distribution Models in Ecology. *Ecological modelling*, 135, 147–186.

Hamada, M. S., Wilson, A., Reese, C. S. & Martz, H. (2008). *Bayesian Reliability*. Springer Science & Business Media.

Harwood, A. J., Metcalfe, N. B., Griffiths, S. W. & Armstrong, J. D. (2002). Intra- and Inter-Specific Competition for Winter Concealment Habitat in Juvenile Salmonids. *Canadian Journal of Fisheries and Aquatic Sciences*, 59, 1515–1523.

Hastings, A. & Wysham, D. B. (2010). Regime shifts in ecological systems can occur with no warning. *Ecology Letters*, 13, 464–472.

Heggberget, T. G. (1988). Timing of Spawning in Norwegian Atlantic Salmon (*Salmo salar*). *Canadian Journal of FIsheries and Aquatic Sciences*, 45, 845–849.

Hines, A. H., Johnson, E. G., Darnell, M. Z., Rittschof, D., Miller, T. J., Bauer, L. J., Rodgers, P. & Aguilar, R. (2010). Predicting Effects of Climate Change on Blue Crabs in Chesapeake Bay. In: *Biology and Management of Exploited Crab Populations under Climate Change*. Chap. Predicting Effects of Climate Change on Blue Crabs in Chesapeake Bay, pp. 109–127.

Ieva, F., Jackson, C. H. & Sharples, L. D. (2015). Multi-state Modelling of Repeated Hospitalisation and Death in Patients with Heart Failure: The Use of Large Administrative Databases in Clinical Epidemiology. *Statistical Methods in Medical Research*,

Jackson, C. (2011). Multi-State Models for Panel Data: The **msm** Package for **R**. *Journal of Statistical Software*, 38.

Jensen, A. J. & Johnsen, B. O. (1999). The Functional Relationship Between Peak Spring Floods and Survival and Growth of Juvenile Atlantic Salmon (*Salmo salar*) and Brown Trout (*Salmo trutta*). *Functional ecology*, 13, 778–785.

Johnsson, J. I., Rydeborg, A. & Sundström, L. F. (2004). Predation Risk and the Territory Value of Cover: An Experimental Study. *Behavioral Ecology and Sociobiology*, 56, 388–392.

Jolly, G. M. (1965). Explicit Estimates from Capture-Recapture Data with Both Death and Immigration-Stochastic Model. *Biometrika*, 52, 225–248.

Jonsson, B. & Jonsson, N. (2009). A Review of the Likely Effects of Climate Change on Anadromous Atlantic salmon *Salmo salar* and Brown Trout *Salmo trutta*, with Particular Reference to Water Temperature and Flow. *Journal of Fish Biology*, 75, 2381–2447.

Kanno, Y., Vokoun, J. C. & Letcher, B. H. (2013). Paired Stream-Air Temperature Measurements Reveal Fine-Scale Thermal Heterogeneity Within Headwater Brook Trout Stream Networks. *River Research and Applications*,

Kendall, W. L., Nichols, J. D. & Hines, J. E. (1997). Estimating temporary emigration using capture-recapture data with Pollock's robust design. *Ecology*, 78, 563–578.

Kendall, W. L. & Nichols, J. D. (2002). Estimating State-Transition Probabilities for Unobservable States Using Capture-Recapture/Resighting Data. *Ecology*, 83, 3276–3284.

Kéry, M. & Schaub, M. (2012). *Bayesian Population Analysis Using **WinBUGS***. Academic Press.

Krkosek, M. & Drake, J. M. (2014). On Signals of Phase Transitions in Salmon Population Dynamics. *Proceedings of the Royal Society of London B: Biological Sciences*, 281.

Lawson, C. R., Vindenes, Y., Bailey, L. & van de Pol, M. (2015). Environmental Variation and Population Responses to Global Change. *Ecology Letters*, 724–736.

Leaniz, C. G. de, Fraser, N. & Huntingford, F. A. (2000). Variability in Performance in Wild Atlantic Salmon, *Salmo salar* L., Fry from a Single Redd. *Fisheries Management and Ecology*, 7, 489–502.

Lebreton, J.-D., Burnham, K. P., Clobert, J. & Anderson, D. R. (1992). Modeling Survival and Testing Biological Hypotheses Using Narked Animals: a Unified Approach with Case Studies. *Ecological monographs*, 62, 67–118.

Letcher, B. H., Dubreuil, T., O'Donnell, M. J., Obedzinski, M., Griswold, K. & Nislow, K. H. (2004). Long-Term Consequences of Variation in Timing and Manner of Fry Introduction on Juvenile Atlantic Salmon (*Salmo salar*) Growth, Survival, and Life-History Expression. *Canadian Journal of Fisheries and Aquatic Sciences*, 61, 2288–2301.

Letcher, B. H. & Horton, G. E. (2008). Seasonal Variation in Size-Dependent Survival of Juvenile Atlantic Salmon (*Salmo salar*): Performance of Multistate Capture-Mark-Recapture Models. *Canadian Journal of Fisheries and Aquatic Sciences*, 65, 1649–1666.

Letcher, B. H., Gries, G. & Juanes, F. (2002). Survival of Stream-Dwelling Atlantic Salmon: Effects of Life History Variation, Season, and Age. *Transaction of the American Fisheries Society*, 131, 838–854.

Letcher, B. H., Schueller, P., Bassar, R. D., Nislow, K. H., Coombs, J. A., Sakrejda, K., Morrissey, M., Sigourney, D. B., Whiteley, A. R. & O'Donnell, M. J. (2014). Robust Estimates of Environmental Effects on Population Vital Rates: an Integrated Capture-Recapture Model of Seasonal Brook Trout Growth, Survival and Movement in a Stream Network. *Journal of Animal Ecology*,

Letcher, B., Walker, J. & Jennison, C. *Spatial Hydro-Ecological Decision System*. URL: `http://ecosheds.org/home`.

Linnansaari, T., Alfredsen, K., Stickler, M., Arnekleiv, J. V., Harby, A. & Cunjak, R. A. (2009). Does Ice Matter? Site Fidelity and Movements by Atlantic Salmon (*Salmo salar* L.) Parr During Winter in a Substrate Enhanced River Reach. *River Research and Applications*, 25, 773–787.

Logez, M., Bady, P. & Pont, D. (2012). Modelling the Habitat Requirement of Riverine Fish Species at the European Scale: Sensitivity to Temperature and Precipitation and Associated Uncertainty. *Ecology of Freshwater Fish*, 21, 266–282.

Lunn, D., Thomas, A., Best, N. & Spiegelhalter, D. (2000). WinBUGS - A Bayesian Modelling Framework: Concepts, Structure, and Extensibility, 10, 325–337.

Lynch, H. J., Naveen, R., Trathan, P. N. & Fagan, W. F. (2012). Spatially Integrated Assessment Reveals Widespread Changes in Penguin Populations on the Antarctic Peninsula. *Ecology*, 93, 1367–1377.

Lyons, J., Stewart, J. S. & Mitro, M. (2010). Predicted Effects of Climate Warming on the Distribution of 50 Stream Fishes in Wisconsin, USA. *Journal of Fish Biology*, 77, 1867–1898.

Metcalfe, N., Huntingford, F. A. & Thorpe, J. (1988). Feeding Intensity, Growth Rates, and the Establishment of Life-History Patterns in Juvenile Atlantic Salmon *Salmo salar*. *The Journal of Animal Ecology*, 463–474.

Metcalfe, N. B. & Thorpe, J. E. (1992). Anorexia and Defended Energy Levels in Over-Wintering Juvenile Salmon. *Journal of animal ecology*, 61, 175–181.

Muenchow, G. (1986). Ecological Use of Failure Time Analysis. *Ecology*, 67, 246–250.

Parrish, D. L., Behnke, R. J., Gephard, S. R., McCormick, S. D. & Reeves, G. H. (1998). Why Aren't There More Atlantic Salmon (*Salmo salar*)? *Canadian Journal of Fisheries and Aquatic Sciences*, 55, 281–287.

Peake, S., McKinley, R. S. & Scruton, D. A. (1997). Swimming Preformance of Various Freshwater Newfoundland Salmonids Relative to Habitat Selection and Fishway Design. *Journal of Fish Biology*, 51, 710–723.

Peterson, A. T., Ortega-Huerta, M. A., Bartley, J., Sanchez-Cordero, V., Soberon, J., Buddemeier, R. H. & Stockwell, D. R. B. (2002). Future Projections for Mexican Faunas Under Global Climate Change Scenarios. *Nature*, 416, 626–629.

Petersson, E., Jarvi, T., Steffner, N. G. & Ragnarsson, B. (1996). The Effect of Domestication on Some Life History Traits of Sea Trout and Atlantic Salmon. *Journal of Fish Biology*, 48, 776–791.

Plummer, M. (2003). JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling. In: *Proceedings of the 3rd international workshop on distributed statistical computing*. Vol. 124. Technische Universit at Wien, pp. 125–132.

Pollock, K. H., Winterstein, S. R., Bunck, C. M. & Curtis, P. D. (1989). Survival Analysis in Telemetry Studies: The Staggered Entry Design. *The Journal of Wildlife Management*, 53, 7–15.

Putter, H., Fiocco, M. & Geskus, R. (2007). Tutorial in Biostatistics: Competing Risks and Multi-State Models. *Statistics in Medicine*, 26, 2389–2430.

Randin, C. F., Dirnböck, T., Dullinger, S., Zimmermann, N. E., Zappa, M. & Guisan, A. (2006). Are Niche-Based Species Distribution Models Transferable in Space? *Journal of Biogeography*, 33, 1689–1703.

Rimmer, D. M., Paim, U. & Saunders, R. L. (1983). Autumnal Habitat Shift of Juvenile Atlantic Salmon (*Salmo salar*) in a Small River. *Canadian Journal of Fisheries and Aquatic Sciences*, 40, 671–680.

Rimmer, D. M., Paim, U. & Saunders, R. L. (1984). Changes in the Selection of Microhabitat by Juvenile Atlantic Salmon (*Salmo salar*) at the Summer-Autumn Transition in a Small River. *Canadian Journal of Fisheries and Aquatic Sciences*, 41, 469–475.

Royle, A. J. & Dorazio, R. M. (2008). *Hierarchical Modeling and Inference in Ecology: The Analysis of Data from Populations, Metapopulations and Communities*. Academic Press.

Rubin, J. F. & Glimsater, C. (1996). Egg-to-Fry Survival of the Sea Trout in Some Streams of Gotland. *Journal of Fish Biology*, 48, 585–606.

Sakrejda, K. *Code Repository*. URL: http://www.fawkes.io.

Saunders, R. L. & Gee, J. H. (1964). Movements of Young Atlantic Salmon in a Small Stream. *Journal of the Fisheries Research Board of Canada*, 21, 27–36.

Scace, J. G., Letcher, B. H. & Noreika, J. (2007). An Efficient Smolt Trap for Sandy and Debris-Laden Streams. *North American Journal of Fisheries Management*, 27, 1276–1286.

Schaub, M. & Abadi, F. (2011). Integrated Ppulation Models: a Novel Analysis Framework for Deeper Insights into Population Dynamics. *Journal of Ornithology*, 152, 227–237.

Schaub, M., Gimenez, O., Sierro, A. & Arlettaz, R. (2007). Use of Integrated Modeling to Enhance Estimates of Population Dynamics Obtained from Limited Data. *Conservation Biology*, 21, 945–955.

Seber, G. A. F. (1965). A Note on Multiple-Recapture Census. *Biometrika*, 52, 249–259.

Sigourney, D. B. (2010). Growth of Atlantic Salmon (*Salmo salar*) in Freshwater". PhD thesis. University of Massachusetts, Amherst.

Snickars, M., Gullström, M., Sundblad, G., Bergström, U., Downie, A.-L., Lindegarth, M. & Mattila, J. (2014). Species-Environment Relationships and Potential for Distribution Modelling in Coastal Waters. *Journal of Sea Research*, 85, 116–125.

Stan Development Team (2014). *Stan: A C++ Library for Probability and Sampling, Version 2.6.2*.

Steingrímsson, S. Ó. & Grant, J. W. A. (2003). Patterns and Correlates of Movement and Site Fidelity in Individually Tagged Young-of-the-Year Atlantic Salmon (*Salmo salar*). *Canadian Journal of Fisheries and Aquatic Sciences*, 60, 193–202.

Sundblad, G., Härmä, M., Lappalainen, A., Urho, L. & Bergström, U. (2009). Transferability of Predictive Fish Distribution Models in Two Coastal Systems. *Estuarine, Coastal and Shelf Science*, 83, 90–96.

Torgersen, C. E., Price, D. M., Li, H. W. & McIntosh, B. A. (1999). Multiscale Thermal Refugia and Stream Habitat Associations of Chinook Salmon in Northeastern Oregon. *Ecological Applications*, 9, 301–319.

Walsh, D. P., Dreitz, V. J. & Heisey, D. M. (2015). Integrated Survival Analysis Using an Event-Time Approach in a Bayesian Framework. *Ecology and Evolution*, 5, 769–780.

Webb, J. H., Fryer, R. J., Taggart, J. B., Thompson, C. E. & Youngson, A. F. (2001). Disperson of Atlantic Salmon (*Salmo salar*) Fry From Competing Families are Revealed by DNA Profiling. *Canadian Journal of Fisheries and Aquatic Sciences*, 58, 2386–2395.

Wenger, S. J. & Olden, J. D. (2012). Assessing Transferability of Ecological Models: an Underappreciated Aspect of Statistical Validation. *Methods in Ecology and Evolution*, 3, 260–267.

White, G. C. & Burnham, K. P. (1999). Program **MARK**: Survival Estimation from Populations of Marked Animals. *Bird study*, 46, S120–S139.

Wolkewitz, M., Allignol, A., Schumacher, M. & Beyersmann, J. (2010). Two Pitfalls in Survival Analysis of Time-Dependent Exposure: A Case Study in a Cohort of Oscar Nominees. *The American Statistician*, 64, 205–211.

Wood, S. N. (2006). Low-Rank Scale-Invariant Tensor Product Smooths for Generalized Additive Mixed Models. *Biometrics*, 62, 1025–1036.

Wreede, L. C. de, Fiocco, M. & Putter, H. (2011). **mstate**: An R Package for the Analysis of Competing Risks and Multi-State Models. *Journal of Statistical Software*, 38, 1–30.

Xenopoulos, M. A., Lodge, D. M., Alcamo, J., Marker, M., Schulze, K. & Van Vuuren, D. P. (2005). Scenarious of Freshwater Fish Extinctions from Climate Change and Water Withdrawal. *Global Change Biology*, 11, 1557–1564.

Yakovlev, G., Rundle, J. B., Shcherbakov, R. & Turcotte, D. L. (2008). *Inter-Arrival Time Distribution for the Non-Homogeneous Poisson Process*. English.

Zydlewski, G. B., Horton, G., Dubreuil, T., Letcher, B., Casey, S. & Zydlewski, J. (2006). Remote Monitoring of Fish in Small Streams: A Unified Approach Using PIT Tags. *Fisheries Research*, 31, 492–502.