Joint Forecasting of Salmon Lice and Treatment Interventions in Aquaculture Operations

BY Benjamin S. Narum and Geir D. Berentsen

DISCUSSION PAPER

Institutt for foretaksøkonomi
Department of Business and Management Science

FOR 7/2024

ISSN: 2387-3000
May 2024

Joint Forecasting of Salmon Lice and Treatment Interventions in Aquaculture Operations

Benjamin S. Narum[1] and Geir D. Berentsen[1]

[1]Department of Business and Management Science, NHH Norwegian School of Economics

23rd May 2024

Abstract

The need for joint forecasting of parasitic lice and associated preventative treatments stems from large monetary losses associated with such treatments, and the distribution of potential future treatments can be used in operational planning to hedge their associated risk. We present a spatio-temporal forecasting model that accounts for the joint dynamics between lice and treatments where spatial interaction between sites is derived from hydrodynamic transportation patterns. The model-derived forecasting distributions exhibit large heterogeneity between sites at significant levels of exposure which suggests the forecasting model can provide great value in assisting operational risk management.

Keywords: Long-term forecasting, GARMA models, density forecasts, aquaculture, salmon lice

1 Introduction

Salmonid aquaculture consists of rearing fish in cages in the ocean, located at sites along the coastline. The immediate environment of the fish is a significant source of uncertainty and parasitic salmon lice, which spread between sites, have been a particular concern for the industry. Salmon lice are a source of biological risk, and their added annual cost is estimated to be in the order of 9% of the annual revenue industry-wide in Norway (Abolofia et al., 2017; Iversen et al., 2017). Farmers manage a portfolio of sites and their overall risk exposure is considerable. In their annual report, the current largest Salmonid farmer world-wide states salmon lice to be one of their major operational risks factors (Mowi ASA, 2023).

We propose that biological risk associated with parasitic lice development can be mitigated through improved operational planning by incorporating spatial diversification between sites. The contribution of this paper is to provide a forecasting model to assess the overall exposure to lice-induced risk as well as heterogeneity between sites. This allows bringing real-time consideration of biological risk into planning. The primary mechanisms for managing biological risk are scheduling of harvest plans (timing and sequencing between sites), as well as reactive treatment interventions against lice. We propose a distributional forecasting model to infer the long-term biological risk that stems from salmon lice as well as their associated treatment interventions. Lice dynamics is a complicated spatio-temporal phenomenon affected by treatment interventions, which themselves are triggered by increased lice abundance. Since sites are connected by company-wide limitations, harvest planning is a complicated large-scale scheduling problem under uncertainty. Finding effective hedging strategies for harvest plans is addressed thorough multistage stochastic programming (King & Wallace, 2012; Shapiro et al., 2014) in a separate publication by applying the forecasting model proposed in this paper.

The salmon louse is a parasite that lives off Salmonids (salmon and trout) and spreads between hosts in the ocean waters. Large expansion of the Norwegian Salmonid aquaculture industry has increased the density of hosts for the salmon louse which, in turn, has resulted in higher levels of infection. This causes concern about the welfare of both farmed and wild Salmonids. As a mitigation strategy, farmers apply removal treatments when they observe high levels of lice abundance, and such treatments contribute significantly to increased fish mortality (Bang Jensen et al., 2020; Kristoffersen et al., 2018; Walde et al., 2021) and loss in growth (Walde et al., 2022). While higher levels of lice abundance is the reason for performing treatments, most of the costs and biological risks are directly associated to the treatments

themselves. Since the aim is operational risk management, our target forecast quantity is the *number of future treatments*. At the levels of lice abundance normally found at aquaculture sites, the biological risk from the parasite itself is comparably low due to the high frequency of treatments. Knowing the exposure to future lice treatments is important for farmers to manage the associated biological risk.

We model treatments as a stochastic process that interacts with another stochastic process for lice abundance. While treatment is a decision made by farmers, there are strong limitations in the flexibility of this decision. We argue that the prospective frequency of future treatments is mainly driven by its interaction with lice abundance, and only moderately by the specific timing that farmers decide, so that it can be modelled as a stochastic process. First, farmers are obliged to treat by regulation based on measured levels of lice abundance, and the time of treatment must lie within a short time-range (± 1 week) once higher levels of lice abundance have been observed. Second, we primarily aim to forecast the long-term treatment count (in the order of 20 weeks ahead), not necessarily the very next treatment, which suggests the flexibility in treatment decisions has a relatively low impact. If short-term plans for treatment (1 week ahead) is provided, we may also simply insert these as explanatory variables before making longer-term forecasts. Lastly, modelling treatments as a stochastic process is also motivated by the fact that these would introduce decision-dependent uncertainty where decisions affect the distribution (Hellemo et al., 2018; Jonsbråten et al., 1998). Since harvest planning models would already be challenging to solve, and since the flexibility in treatment decisions is considered to be low, decision-dependent uncertainty is avoided by making treatment decisions exogenous with respect to harvest decisions.

Existing studies on the dynamics and spread of lice can broadly be divided into two categories. In the first category are statistical models used to draw inference about the dynamics of lice or to obtain site-specific short-term forecasts (Aldrin et al., 2017; Aldrin et al., 2019; Aldrin et al., 2013; Elghafghuf et al., 2020; Elghafghuf et al., 2018). Some of these models also incorporate spatial dependence by considering seaway distance (along the ocean surface) between sites. In the second category is a physical model that uses simulation of hydrodynamic stream patterns combined with historical lice abundance data to infer a measure of *infestation pressure* (Myksvoll et al., 2018). A caveat of this latter approach is that it cannot be used forecast future lice counts since the relation between infestation pressure and on-site dynamics of lice is not accounted for. To apply a forecasting model to operational planning, spatial phenomena must be truthfully reflected together with on-site dynamics. Our approach bridges the gap between the preexisting two approaches by introducing asymmetric spatial effects derived from stream simulation data (as opposed to the symmetric seaway distance) in a time-series model that links infestation pressure to the on-site dynamics of lice. Moreover, while existing approaches consider lice treatments as an explanatory variable, we consider lice and treatments as two stochastic processes that affect each other. To our knowledge, this paper is the first to incorporate stream patterns into on-site lice dynamics as well as its interaction with future treatments. Furthermore, previous literature does not emphasize longer forecasting horizons, nor distributional forecasting, as required for forward-looking harvest planning under uncertainty.

The paper is structured as follows: Section 2 presents all data sources and explains the basics of lice development; Section 3 presents the forecasting models for lice and treatments, with details on joint forecasting and estimation; Section 4 presents results and model validations; finally, Section 5 is a discussion with concluding remarks.

2 Data

We limit sites to those in the Norwegian production area 3 situated on the western coast of Norway, which is the largest production area known for having the most activity of lice. Sites outside of this area are incorporated by their contribution to neighbouring infection. There are 137 sites in total that are active for 41572 site-weeks combined within the period June 2012 to September 2021. An overview of these sites is illustrated in Figure 1.

2.1 Lice and treatment data

Biologically, lice develop through multiple stages (Hamre et al., 2013) that we simplify into categories relevant to our model (see Figure 2). The development cycle starts by hatching larvae (LR) from egg strings attached to adult female (AF) lice. The larvae are then transported by ocean currents and may find a new host (a Salmonid). During transportation the louse may attach to a new host, then referred to as a *recruit* (R). If it cannot find a new host in time, it runs out of nutrition and dies. Note that the recruit may also originate from the same site. The recruit develops through three stages on the

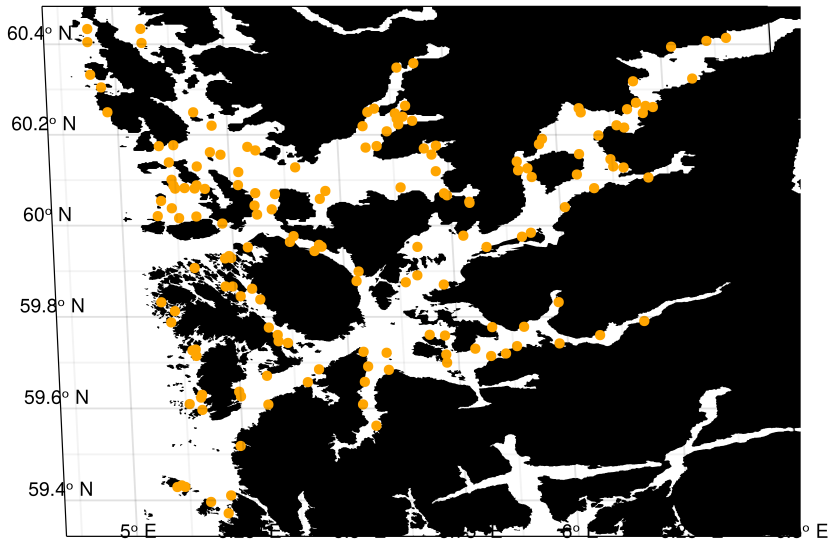

Figure 1: Overview of aquaculture sites in production area 3.

host: stationary (ST); pre-adult (PA), where it starts to move on the host; and lastly, adult female/male (AF/AM), where the female develops egg strings to hatch new larval lice (LR) that repeat the cycle. Keep in mind that this development cycle is highly temperature dependent and that males develop slightly faster than female lice (Hamre et al., 2019). Also note that pre-adult (PA) and adult male (AM) lice are visually indistinguishable and are thus only identified as mobile lice (MB).

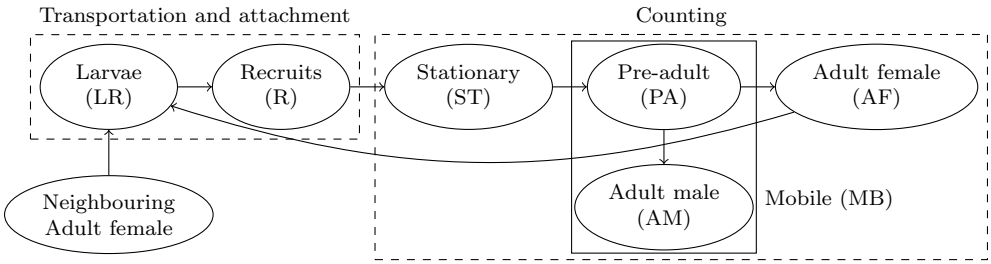

Figure 2: Different stages of the lice development cycle used for the present model. Transportation happens in the two first stages (LR and R), and lice are counted in the latter three (ST, MB = PA + AM and AF) while attached to the fish.

By regulation, all Norwegian farmers are obliged to count lice in the three identifiable categories (ST, MB and AF) weekly on 10 or 20 fish (depending on the time of year and location) in each production cage (Lovdata, 2012). Typically, a single cage has approximately 200 000 fish in total and the lice count on each fish sample is averaged over all cages at a site (1–12 cages per site). This average is referred to as the *lice abundance* at a given site. These weekly reports are downloaded through Barentswatch[1]. Missing lice counts are replaced by interpolation between the nearest foregoing and upcoming reports, reduced by 70% per week from the interpolation points (1597 interpolated points in total).

Farmers are obliged to commence lice treatment on the entire site when lice abundance is measured above 0.2 or 0.5 (depending on the time of year and location) (Lovdata, 2012). The lower limit was changed from 0.1 to 0.2 in 2017 (Lovdata, 2017). The weekly lice reports collected from Barentswatch also contain information on treatments, and these mainly distinguish between medical, mechanical, and cleaner fish treatments. Cleaner fish are other species of smaller fish deployed in the cage who eat lice off the Salmonids. Partial treatment of the site is allowed only by particularly good reason; hence, we do not distinguish between partial and complete treatments. Treatments are reported weekly but may last for more than one week. For medical treatments, farmers first report the start date and later report the end date of the treatment. We use the starting week as the occurrence of medical treatments.

[1] https://www.barentswatch.no/fiskehelse/

Medical treatments with missing start dates are placed at the week before the reported end of treatment. Treatments of the same kind (medical or mechanical) reported multiple times in consecutive weeks are joined together as a single treatment and placed on the first week of reporting. This prevents duplicates of the same treatment taking place over several weeks. If there are more than a certain number of treatments in consecutive weeks (3 for mechanical and 4 for medical), these are split as separate treatment periods, with commencement at the first week of each period. By visual inspection, we find that some sites have missing treatment reports. These sites are removed when fitting the models (4 sites in total), but still contribute to neighbour effects and are included in simulations. There are still a few minor active time periods where we suspect treatment reports are missing but these are still included. Cleaner fish treatments are not necessarily reported after 2018 (Lovdata, 2018); for this reason, the effect of cleaner fish is not included when forecasting treatments but is included as a corrective term during model estimation. Figure 3 illustrates the number of treatments in production area 3 within the relevant time range, and we see a clear shift from medical treatment to mechanical treatment. This is due to the development of resistance to commonly used medical substances, which caused a transition to mechanical treatments instead (Jensen et al., 2020). There are treatments in 8.4% of active weeks (3499 in total).

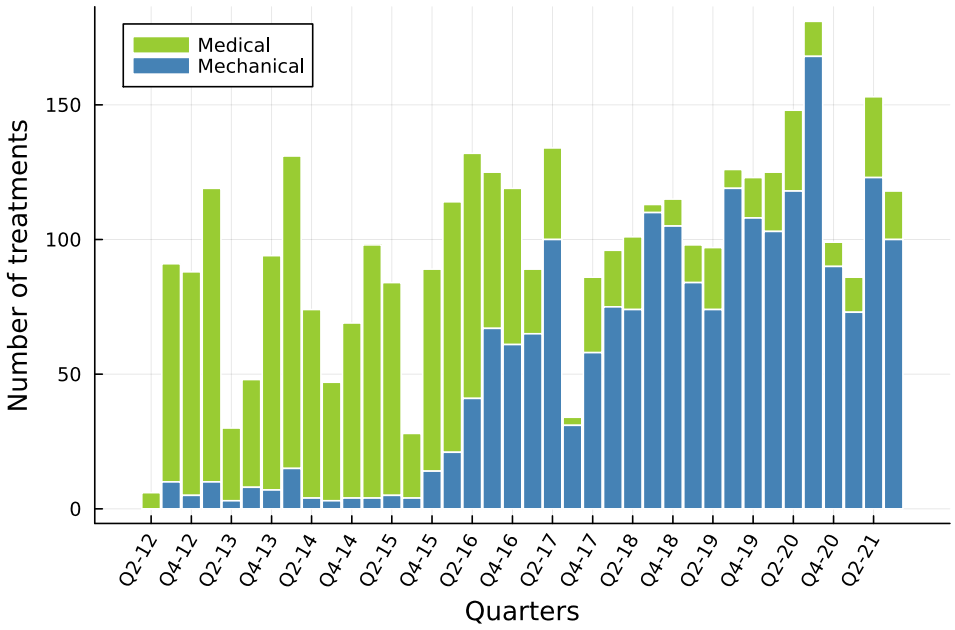

Figure 3: Treatments in production area 3.

Salmonid farming in Norway requires two months of fallowing (closing the site) at the end of each production cycle as a preventive measure against lice, and there are no lice reports during this time (Lovdata, 2008). While a site is active, farmers are obliged to report lice every week, and we can know which sites are active based on the presence of reports. However, some reports are missing in the data. We must distinguish between inactivity and missing reports since inactivity means lice abundance is zero while missing reports do not. If there are more than six weeks of consecutively missing reports, we assume the site is inactive, while otherwise, we assume reports are missing and interpolate the missing data. This is conservative by letting some (shorter) inactive periods potentially seem active instead, which prevents zeroing out the dynamics of lice abundance. We assume there are no missing treatments in weeks of missing reports.

Several sources of uncertainty are related to the lice count and treatment data. An extensive survey on this was conducted and summarized by Solberg et al. (2018). First, lice counts are known to have considerable measurement error, mainly attributed to the small sample of fish collected to count lice. The small sample is motivated by negative health implications from handling the fish when counting lice. Second, there is uncertainty about whether the human counter can identify lice on the fish and correctly classify the stage of the louse. In particular, lice in the stationary stage (ST) are especially small in size and can be hard to observe (Thorvaldsen et al., 2019). Third, there are doubts as to whether the sample of fish collected for counting is representative of all fish in a cage since these are collected only from the sea surface. Fourth, regulation allows counting lice multiple times in each cage and reporting only the last counting, which enables selective re-counting. Fifth, due to the high cost of treatments and the

low treatment limit, farmers have clear incentives to under-report lice abundance and avoid treatments, which can affect the quality of the data. The impact of such incentives on this dataset was explored by Jeong et al. (2023). Lastly, regulation allows treatments to be dismissed for the last three weeks before fallowing, which can explain some higher lice counts at the end of production cycles. Recently, technology to count lice by camera technology has been developed and is already in use at some sites. This means more reliable count data will become available over time.

2.2 Temperature and transportation by coastal currents

Temperature estimates at three-meter depth are collected from the NorKyst800 model (Albretsen et al., 2011), which describes coastal conditions in the form of currents and temperature on a 800m by 800m grid. This model has been validated to give high-quality temperature estimates with a deviation of at most $1°C$ (Asplin et al., 2020). Occasional missing data within short time spans are filled by linear interpolation. We assume temperature forecasts can be obtained from other models (for example, from the European Centre for Medium-Range Weather Forecasts[2]) and consider it to be a given explanatory variable.

Transportation of lice between site pairs is mainly governed by stream patterns. We derive transportation patterns from a particle simulation model developed by the Norwegian Institute of Marine Research (see Myksvoll et al., 2018, for details), which uses stream patterns from the NorKyst800 model to simulate how particles drift. The particle simulation model is currently used in a national monitoring system to regulate the industry and has been validated against field experiment data to have high accuracy (Myksvoll et al., 2018). In detail, the simulation runs as follows: For every hour, three particles are released from every site and tracked as they follow stream patterns determined by NorKyst800. Whenever a particle is in proximity to another site, it is registered as an encounter with information about the source site, destination site, and age in hours. Particles whose age in degree-days surpasses 200 are removed. The simulation runs for nine months of stream data from February to October 2017. Sites that were not active during this simulation window have been removed (17 sites). From these simulations, we obtain a total transportation volume between pairs of sites, referred to as *connectivity* (in number of particles), as well as an empirical distribution for the transportation time. Both connectivity and transportation time are asymmetric between site pairs, and Figure 4 illustrates the relationship between connectivity and seaway distance used in previous literature (Aldrin et al., 2017; Aldrin et al., 2019; Jansen et al., 2012). Seaway distance is symmetric, and its discrepancy with connectivity can be large.

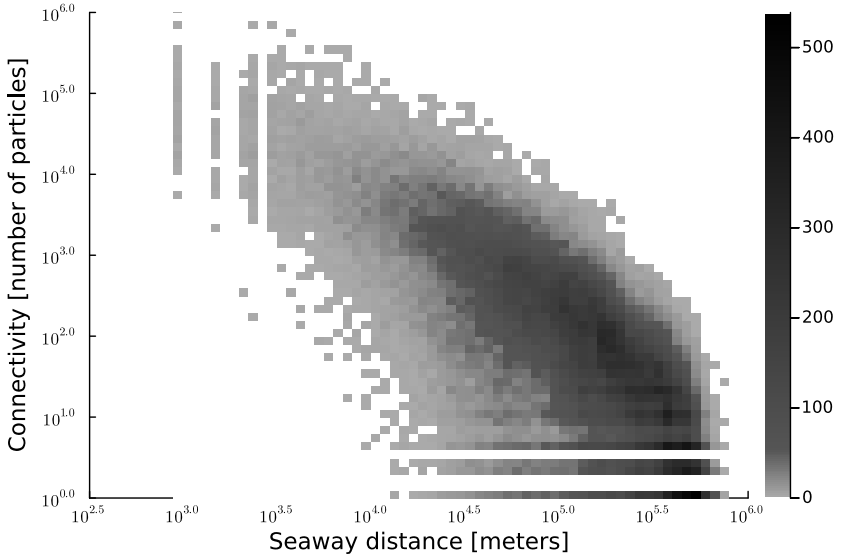

Figure 4: Comparison of seaway distance (in meters) to connectivity (number of particle encounters) among all site pairs.

[2]www.ecmwf.int/

3 Methodology

To develop a model that can forecast treatments, we consider lice abundance and treatments as two processes that affect each other. Successful treatment should lead to lower lice abundance, while higher lice abundance should lead to increased probability of treatment. Furthermore, we assume the occurrence of treatment has a small delay so that it may only depend on past lice abundance; meanwhile, lice abundance may also depend on the occurrence of treatment within the same week. Figure 5 illustrates this lead-lag relationship. This delay assumption on treatments allows constructing a joint forecasting model based on separate models for the conditional marginal distribution of lice abundance and treatment. These models are described in Section 3.1 and Section 3.2, respectively, while joint forecasting using both models is described in Section 3.3. The estimation procedure is described in Section 3.4.

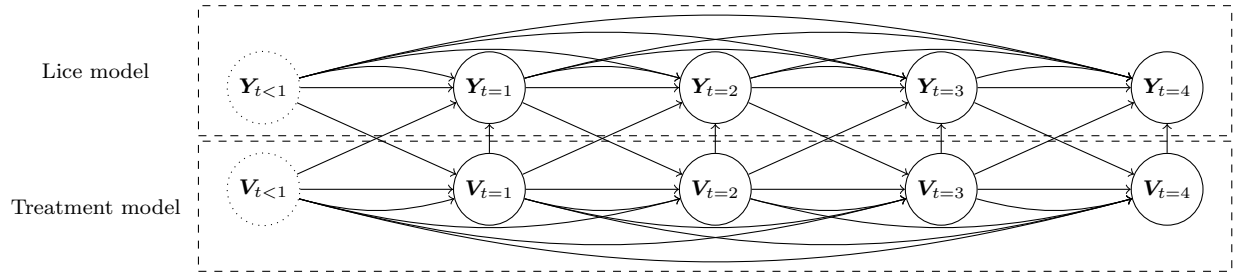

Figure 5: Dynamics between lice ($\boldsymbol{Y}$) and treatments ($\boldsymbol{V}$) at different time steps.

Each of the models are designed in a two-level hierarchy having slightly different interpretations. First, we derive a *best-guess estimate* that incorporates all explanatory variables and domain knowledge about the respective phenomena. Second, we correct any remaining time-structure using Generalised Auto-Regressive Moving-Average (GARMA) models (Benjamin et al., 2003). For the lice model, the best-guess estimate ensures consistent long-term forecasting ability based on explicit knowledge about the development-cycle of lice, while a moving-average component improves short-term predictions by correcting for more recent observations of actual lice abundance. For the treatment model, the best-guess estimate acts as a current assessment of the probability of treatment (mainly based on lice abundance) while an auto-regressive component corrects for relations to past treatments.

3.1 Lice abundance model

Let $n_{it}Y_{it}^{\star}$ denote the observed *lice count* within stages $\star \in \{\mathrm{ST}, \mathrm{MB}, \mathrm{AF}\}$ on n_{it} fish at site $i \in \{1, \ldots, N\}$ in week t. Correspondingly, $Y_{it}^{\star}$ denotes the observed *lice abundance*; namely, the average number of lice per fish. We represent the *occurrence of treatment* by a random variables $V_{it}^{\diamond}$ where $\diamond \in \{\mathrm{mec}, \mathrm{med}, \mathrm{clf}\}$ denotes the different kinds of treatments: mechanical, medical or cleaner fish. Cleaner fish V_{it}^{clf} is a continuous rate, representing the number of cleaner fish deployed and is scaled by the capacity of the site. The binary variables $V_{it}^{\diamond}$ for $\diamond \in \{\mathrm{mec}, \mathrm{med}\}$ take value one if treatment $\diamond$ is commenced at site i in week t, and zero otherwise. We let $\mathcal{F}_t$ denote the information set generated by $Y_{is}^{\star}$ and $V_{is}^{\diamond}$ for all sites $i = 1, \ldots, N$, weeks $s = 1, \ldots, t$, stages $\star \in \{\mathrm{AF}, \mathrm{MB}, \mathrm{ST}\}$, and treatment kinds $\diamond \in \{\mathrm{mec}, \mathrm{med}, \mathrm{clf}\}$. For simplicity, we also let $V_{it} = \left(V_{it}^{\mathrm{mec}}, V_{it}^{\mathrm{med}}, V_{it}^{\mathrm{clf}}\right)$ denote the concatenation of all treatment kinds.

The response variables in the lice model are the lice counts $n_{it}Y_{it}^{\star}$ for all stages $\star$ which we, conditional on $\mathcal{F}_{t-1}$ and V_{it}, assume follow a negative binomial distribution

$$\left(n_{it}Y_{it}^{\star} \mid \mathcal{F}_{t-1}, V_{it}\right) \sim \mathrm{NegBin}\left(n_{it}\mu_{it}^{\star}, n_{it}\nu^{\star}\right), \tag{1}$$

parameterised by its expectation $n_{it}\mu_{it}^{\star}$ and dispersion $n_{it}\nu^{\star}$. We assume a global dispersion parameter $\nu^{\star}$ per stage, but scale by n_{it} to account for increased certainty in the abundance estimate when counting lice on more fish. The negative binomial distribution has been validated experimentally to be well suited for lice counts on farmed salmon (Jeong & Revie, 2020). For the response variable $n_{it}Y_{it}^{\star}$, we have $n_{it} \in [10, 140]$ and $Y_{it}^{\star}$ typically in the range $[0, 1]$.

The dynamics of lice abundance is modelled via the expected abundance $\mu_{it}^{\star}$. As an baseline explanatory term, we use a best-guess of lice abundance $\bar{\mu}_{it}^{\star}$ derived from past recruits and their development into each stage $\star$. While the best-guess estimate has strong physical motivation that improves long-term

forecasting ability, we also have regular observations of actual lice abundance in each stage that should correspond to $\bar{\mu}^\star_{it}$. Any discrepancy between the anticipated lice abundance $\bar{\mu}^\star_{it}$ and observed abundance $Y^\star_{it}$ is corrected for by a Generalised Moving-Average (GMA) model for $\mu^\star_{it}$ to improve short-term predictions. Specifically, we let

$$\log(c + \mu^\star_{it}) = \log(c + \bar{\mu}^\star_{it}) + \sum_{l=1}^{20} A_{itl}\theta_l \zeta^\star_{i,t-l}, \tag{2}$$

where

$$\zeta^\star_{it} = \log(c + Y^\star_{it}) - \log(c + \mu^\star_{it}), \tag{3}$$

denotes the innovation of $\mu^\star_{it}$ with respect to the observation $Y^\star_{it}$. The fitted parameter $c > 0$ determines the zero-level of observations on the logarithmic scale to prevent values of $-\infty$ whenever $Y^\star_{it}, \bar{\mu}^\star_{it}, \mu^\star_{it} = 0$. By design, we then have that $\zeta^\star_{it} = 0$ whenever $Y^\star_{it} = \mu^\star_{it}$, meaning we default to the best-guess estimate whenever this discrepancy is low. The indicator A_{itl} takes value one if farm i is active from week $t - l$ to week t and zero otherwise; this terminates the memory of the process whenever the site has been closed. Through model estimation, we find that the parameters θ_l have a (close to) exponential decay in their values for increasing lags l; hence, we let

$$\theta_l = \theta^{\text{scale}} \exp(\theta^{\text{rate}}(l - 1)), \tag{4}$$

and estimate only the two parameters θ^{scale} and θ^{rate}. We use the same parameters θ^{scale} and θ^{rate} for this lag structure in all stages $\star$, and this is motivated by the fact that we get very similar estimates if we consider these as separate quantities.

The primary aim of the lice model is to give a precise account of the dynamics of lice abundance while incorporating the effect of treatment, and the best-guess estimate of lice abundance $\bar{\mu}^\star$ is motivated by the stage-structured development cycle of lice. For this purpose, we also consider the *unobserved* abundances $Y^{\text{RN}}, Y^{\text{RS}}$ and Y^{RU} that decomposes all recruits (R) into their respective sources: neighbouring sites (RN), within-site (RS), or unexplained sources (RU). The best-guess estimate then takes expression

$$\bar{\mu}^\star_{it} = \kappa_{it0} \sum_{l=1}^{20} A_{itl}\kappa_{itl} r^\star_{itl} \left(\hat{Y}^{\text{RS}}_{i,t-l} + \hat{Y}^{\text{RN}}_{i,t-l} + \hat{Y}^{\text{RU}}_{i,t-l} \right), \tag{5}$$

which incorporates *estimates* $(\hat{Y}^{\text{RS}}, \hat{Y}^{\text{RN}}, \hat{Y}^{\text{RU}})$ of past recruits, while accounting for survivability (κ) and temperature-dependent development times $(r^\star)$. The effect of treatment is accounted for by survivability, which is explained in Section 3.1.1, while the term $r^\star$ that accounts for development time is described in Section 3.1.2. The estimates $\hat{Y}^{\text{RN}}, \hat{Y}^{\text{RS}}$ and $\hat{Y}^{\text{RU}}$ for unobserved abundances are based on hydrodynamic simulations, biological relations, and past observations of adult female lice; and these are described in Section 3.1.3. We use a maximum lag of 20 weeks since most contributions beyond 20 weeks are small when accounting for both development time and mortality.

3.1.1 Mortality

We let the factor κ_{its} represent the survivability of lice. Specifically, the log-survivability rate at site i between weeks $t - s$ and t takes the expression

$$\log(\kappa_{its}) = s\rho + \sum_{l=0}^{s} \left(V^{\text{mec}}_{i,t-l} \sum_{k=0}^{l} \delta^{\text{mec}}_k + V^{\text{med}}_{i,t-l} \sum_{k=0}^{l} \delta^{\text{med}}_k + V^{\text{clf}}_{i,t-l} \sum_{k=0}^{l} \delta^{\text{clf}}_k \right), \tag{6}$$

where the parameter $\rho < 0$ represents a baseline weekly mortality. The parameters $\delta^\diamond_k < 0$ denote the additional log-effect of various treatments, $k \in \{0, 1, \dots\}$ weeks after they were initiated. The total effect of a single treatment initiated l weeks ago then equals $\exp\left(\sum_{k=0}^{l} \delta^\diamond_k\right)$. Hence, κ_{its} is the accumulated survival rate between weeks $t - s$ and t which incorporates all natural and treatment-induced mortality. Observe also that κ_{it0} incorporates the effect of treatment in the current week but omits the baseline weekly mortality ρ.

3.1.2 Development time

To account for development time, we derive auto-regressive coefficients to infer the abundance of counted lice (Y^{ST}, Y^{MB} and Y^{AF}) from the abundance of recruits (Y^{RS}, Y^{RN} and Y^{RU}) in the past. These auto-regressive coefficients are derived as analytical expressions of the temperature-dependent development time. The advantage of this analytical approach is that we may use existing knowledge about the effect of temperature on development time. This avoids the need for re-estimating these effects, and only the potential variability in development time is estimated.

The development time going from a recruit into each consecutive development stage was examined by Hamre et al. (2019), who found that the development time (in weeks) between stages is independent of age but highly dependent on temperature as well as gender. Let T denote temperature, in general, and T_{it} the temperature at site i in week t. The temperature-dependent development time for male and female lice, respectively, take expressions

$$D_M(T) = \exp\left(1.7216 - 0.2472T + 0.0050T^2\right), \tag{7}$$

$$D_F(T) = \exp\left(1.8033 - 0.2172T + 0.0039T^2\right), \tag{8}$$

which are estimated from the data in (Hamre et al., 2019) (see Supplementary Material). Available data covers temperatures in the range 6–21°C, and we extrapolate outside of these (for 14.7% of active site-weeks). The development time from recruit into stationary, pre-adult and adult lice are multiples $1, 3$ and 5 of $D_M(T)$ and $D_F(T)$, as illustrated in Figure 6. We also assume lice die of old age after $10D_M(T)$ and $10D_F(T)$ weeks.

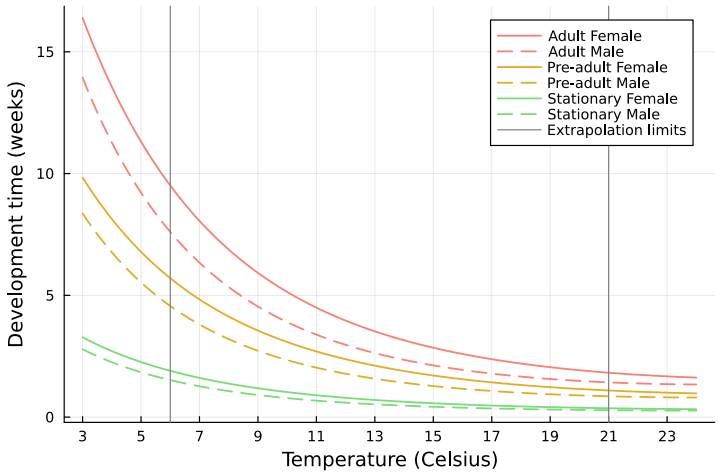

Figure 6: Development time from recruit to the respective stages for both genders at different temperatures.

We use the expressions (7) and (8) to provide estimates for the average development time; however, to account for potential variability in development time, we assume a distribution for development time. We assume development time from recruit into stage $\star$ is Weibull distributed, parameterised by its expectation $\psi_\star$ and shape parameter $\alpha > 0$. Its cumulative distribution function (CDF) is then expressed as

$$F_\alpha(s \mid \psi_\star) = 1 - \exp\left(-\left(\frac{s\Gamma(1+1/\alpha)}{\psi_\star}\right)^\alpha\right), \tag{9}$$

where $\Gamma(\cdot)$ is the gamma function, and the parameter α is subject to estimation. Using this CDF, we may now express the share of recruits that have developed into stage $\star$ (with average development time $\psi_\star > 0$) and *not* into the consecutive stage (with average development time $\psi_{\star+1} \geq \psi_\star$) within s weeks as:

$$F_\alpha(s \mid \psi_\star) - F_\alpha(s \mid \psi_{\star+1}) \geq 0. \tag{10}$$

This quantity is non-negative since these distributions have the same shape parameter α while one expectation is larger than the other. The difference (10) specifies the contribution of recruits s weeks ago to each consecutive stage $\star$ in the current week, and we may sum over all past weeks to find the total contribution to current lice abundance. Due to asymmetries in the contribution of each gender to

different stages, we get the following temperature-dependent coefficients to represent development:

$$r_{its}^{\mathrm{AF}} = F_\alpha(s \mid 5D_F(T_{it})) - F_\alpha(s \mid 10D_F(T_{it})), \tag{11a}$$

$$r_{its}^{\mathrm{MB}} = F_\alpha(s \mid 3D_F(T_{it})) - F_\alpha(s \mid 5D_F(T_{it})) + F_\alpha(s \mid 3D_M(T_{it})) - F_\alpha(s \mid 10D_M(T_{it})), \tag{11b}$$

$$r_{its}^{\mathrm{ST}} = F_\alpha(s \mid 1D_F(T_{it})) - F_\alpha(s \mid 3D_F(T_{it})) + F_\alpha(s \mid 1D_M(T_{it})) - F_\alpha(s \mid 3D_M(T_{it})). \tag{11c}$$

These have the interpretation that $r_{its}^{\star}$ is the rate at which recruits from week $t-s$ contribute to lice abundance in stage $\star$ and week t. The sources contributing to AF are only female lice in the adult stage, while the contribution to MB are pre-adult female lice as well as pre-adult and adult male lice. The contribution to ST is from male and female lice in the stationary stage. Lastly, we also consider mortality from old age for stages MB and AF. The inferred lead-lag relationship between recruits and each stage $\star$ (while ignoring the effect of mortality) is illustrated for different temperatures in Figure 7.

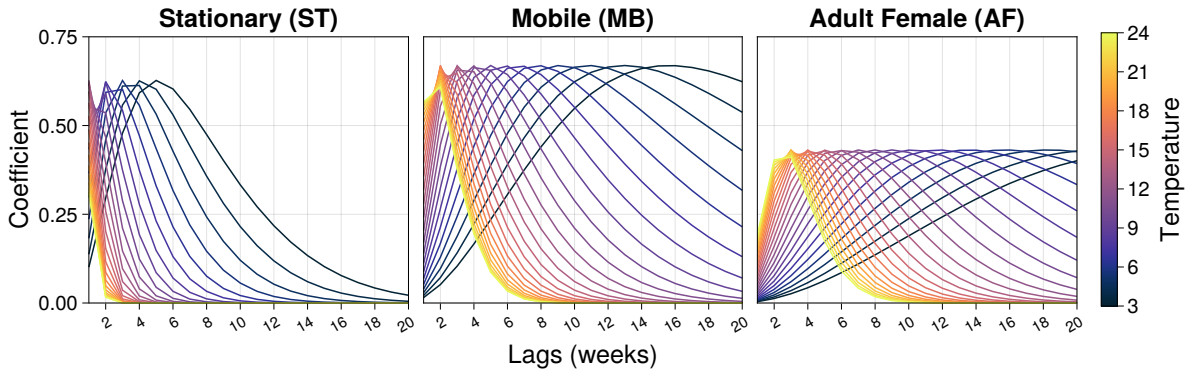

Figure 7: Lead-lag relationship between recruits and the respective stages (ST, MB and AF) that infer temperature-dependent development times.

3.1.3 Estimates of past recruits

Past recruits can be decomposed into different sources and estimated based on past abundance of adult female lice, transportation patterns and successful attachment. In the following, we assume transportation ends when a louse has successfully attached to a new host; then defined to be a recruit. The number of fish at location i, and thereby the total number of lice inferred by lice abundance, are generally not available since it is considered market-sensitive information. Instead, we use the capacity C_{it} of site i in week t (capacity will rarely change over time) as a proxy for fish count. The rate of egg production is temperature-dependent, whose expression is given by

$$H(T) = \exp\left(-0.869324 + 0.15615T - 0.007699T^2\right). \tag{12}$$

This is estimated from data in (Samsing et al., 2016) and normalised to take values close to one (see Supplementary Material).

The estimate of recruits from neighbouring sites then takes expression

$$\hat{Y}_{it}^{\mathrm{RN}} = \frac{\iota^{\mathrm{RN}}}{C_{it}} \sum_{j \in J_i} \sum_{s=1}^{5} w_{jits} \left[C_{j,t-s} H(T_{i,t-s}) Y_{j,t-s}^{\mathrm{AF}} \right], \tag{13}$$

where $C_{j,t-s}H(T_{j,t-s})Y_{j,t-s}^{\mathrm{AF}}$ is a proxy for the number of emitted larvae from site j in week $t-s$, and $\iota^{\mathrm{RN}} > 0$ is a scaling parameter. The weights w_{jits} represent the magnitude of transported lice that originate from site j in week $t-s$ that successfully attaches to a new host at site i in week t. The set J_i contains all relevant neighbours of site i, limited to those that are shown to successfully transport particles to site i in the hydrodynamic simulation. The maximum number of transportation weeks is set to five (see point (iii) below). The weights w_{ijts} are composed of four main factors: (i) total transportation rate as well as the distribution of transportation time, (ii) daily mortality rate before reaching a new host, (iii) time windows to find a new host, and (iv) attachment success rate once reaching a new host. These are modelled as follows (additional details are provided in the Supplementary Material):

(i) Let K_{ji} denote the connectivity from site j to i as number of particles successfully transported in the simulation (in millions). We also derive an empirical probability f_{jis} of transporting particles from site j to i within $s-1$ to s weeks based on transportation times. Overall, $K_{ji}f_{jis}$ is the number of particles transported from site j to i within $s-1$ to s weeks.

(ii) The weekly survival rate is set to $\omega = \exp\left(\frac{7}{4}\log(0.5)\right) = 0.2973$ based on (Myksvoll et al., 2018; Stien et al., 2005), with the interpretation that the survival rate has a half-life of four days.

(iii) There is a specific time window where a successful transmittance can take place. The lower limit of the window comes from the fact that larvae must develop before they attach to a new host. The upper limit comes from the fact that they run out of nutrition and die within a certain number of weeks. This time window is illustrated in Figure 8. Within the relevant temperature range, there is a maximum transportation time of 5 weeks.

(iv) Infectivity is defined as the rate of successful infection once reaching a new host, which is both age and temperature dependent (Skern-Mauritzen et al., 2020). We have estimated a function $I(A,T)$ for infectivity based on data from (Samsing et al., 2016; Skern-Mauritzen et al., 2020) (see Supplementary Material), and use the average infectivity during a week $\bar{I}_s(T)$ to infer infestation success. The function $I(A,T)$ is illustrated in Figure 8.

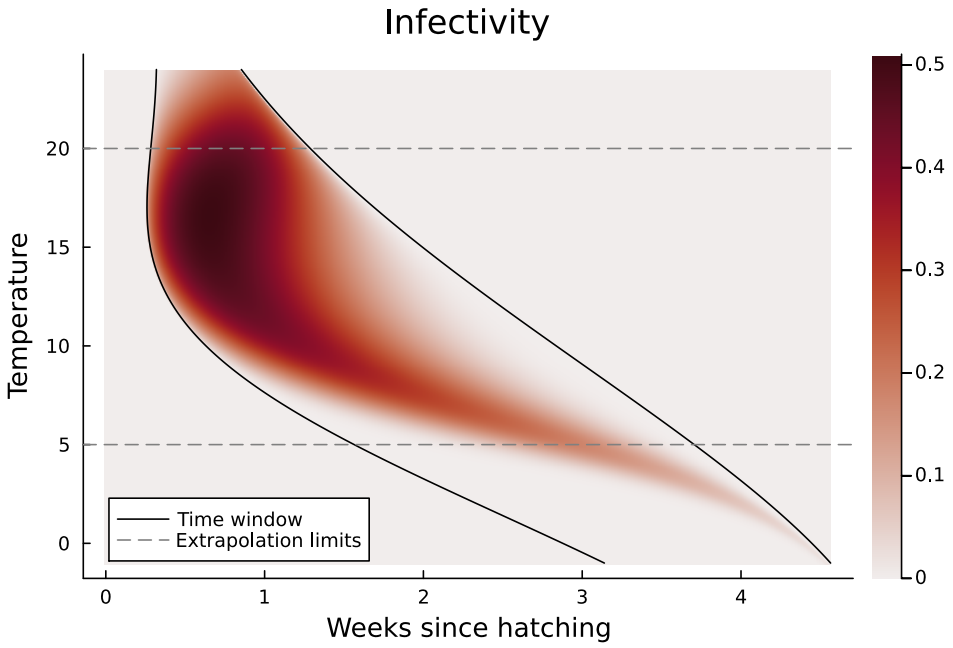

Figure 8: Infectivity $I(A,T)$ as a function of weeks since hatching (A) and temperature (T). The infection window is shown in black. In the model, we extrapolate beyond the data range of temperatures from 5°C to 20°C (dashed lines) to temperatures from −1°C to 24°C (indicated by the axis limits).

Based on (i)–(iv), the weights that account for successful transportation and attachment are given as:

$$w_{jits} = A_{it}A_{j,t-s}\bar{I}_s(T_{it})K_{ji}f_{jis}\omega^s, \quad s \in \{1,\ldots,5\}. \tag{14}$$

The indicator A_{it} is equal to one if site i is active in week t and zero otherwise; hence, $A_{it}A_{j,t-s}$ takes value one if both the emitting site is active during emission and the receiving site is active during arrival. In Figure 9, we illustrate the logarithm of the aggregated transferal and attachment rate between sites (i.e., $\log(\sum_{s=1}^{5} w_{ijts})$ for all i and j) for $T = 9°C$ and assuming all sites are active.

The estimate of within-site recruits $\hat{Y}_{it}^{\mathrm{RS}}$ is derived from the same components as before but excluding travel time. Its estimate takes expression

$$\hat{Y}_{it}^{\mathrm{RS}} = \iota^{\mathrm{RS}}A_{it}\sum_{s=1}^{5} A_{i,t-s}\bar{I}_s(T_{it})H(T_{i,t-s})Y_{i,t-s}^{\mathrm{AF}}, \tag{15}$$

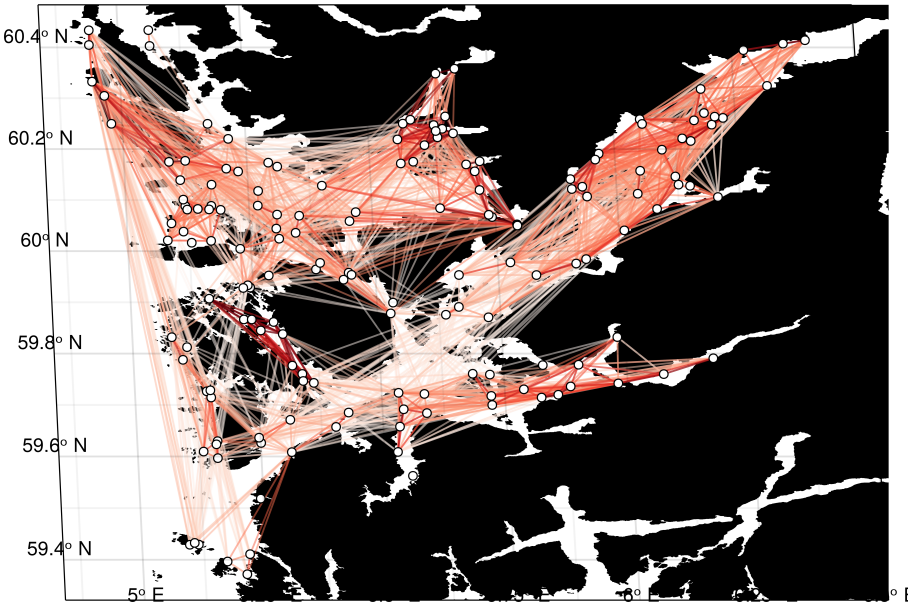

Figure 9: Logarithm of the aggregated transferal and attachment rate between sites at temperature $T = 9°C$ assuming all sites are active, in production area 3. Only rates above the 80th quantile are shown.

which accounts for the past abundance of adult female lice, hatch rate, infectivity, and the active status of the site. The fitted parameter $\iota^{\mathrm{RS}} > 0$ is a scaling.

Lastly, the estimate of recruits from unexplained sources $\hat{Y}^{\mathrm{RU}}$ mainly acts as an intercept term, expressed as

$$\hat{Y}_{it}^{\mathrm{RU}} = \iota^{\mathrm{RU}} \exp(\beta_T (T_{it} - 10)). \tag{16}$$

This is limited to a temperature correction to reflect seasonality and scaling by the parameter $\iota^{\mathrm{RU}} > 0$. By letting the intercept term contribute to recruits instead of $\bar{\mu}^\star$ directly, we get a gradual build-up of lice abundance directly after a site has been re-opened and the intercept will also be affected by treatments.

3.2 Treatment model

When forecasting treatments, we are mainly interested in the the potential need for future treatments rather than the specific kind of treatment; thus, we model whether a treatment occurs or not. Cleaner fish are mainly used as a preventative measure and are ignored for this purpose. Let $W_{it} = V_{it}^{\mathrm{mec}} \wedge V_{it}^{\mathrm{med}}$ denote the binary random variable for whether a mechanical or medical treatment occurs at site i in week t. We then assume the occurrence of treatment follows the conditional distribution

$$(W_{it} \mid \mathcal{F}_{t-1}) \sim \mathrm{Bernoulli}(p_{it}), \tag{17}$$

where p_{it} is the probability of treatment and $\mathcal{F}_{t-1}$ is the information set up to time $t-1$ as defined in Section 3.1. This probability is modelled as a Generalised Auto-Regressive (GAR) model with explanatory variables using a logistic response function.

We express the probability of treatment through $\eta_{it} = \mathrm{logit}(p_{it})$ and let $\bar{\eta}_{it}$ denote the best-guess estimate of the log-odds for treatment based on explanatory variables. This best-guess estimate is given as

$$\bar{\eta}_{it} = \beta_0^p + \beta_L^p L_{it} + \sum_{\star \in \{\mathrm{MB,AF}\}} \left(\gamma_l^\star + \gamma^{I,\star} W_{i,t-1}\right) \left(Y_{i,t-1}^\star\right)^{\pi^\star}, \tag{18}$$

where β_0^p is an intercept, and the indicator variable L_{it} takes value one if the treatment limit is low (i.e., the limit is below 0.5) and zero otherwise. The last term represents the effect of observed lice abundance in the previous week, as well as its interaction with last week's treatment. We find that no further lags on lice abundance are required due to small and insignificant coefficient estimates. Using stationary lice

(Y_{it}^{ST}) to infer treatment has a negligible effect; hence, it is omitted. The parameters $\beta_0^p, \beta_L^p, \gamma_l^\star$, $\gamma^{I,\star}$ and $\pi^\star$ (non-linearity) for $\star \in \{\mathrm{AF}, \mathrm{MB}\}$ are subject to estimation.

We also correct this best-guess estimate ($\bar{\eta}$) using auto-regressive relationships on its past prediction errors to improve its forecasting ability. The corrected log-odds of treatment (η) takes expression

$$\eta_{it} = \bar{\eta}_{it} + \sum_{l=1}^{20} A_{itl}\phi_l \xi_{i,t-1}, \tag{19}$$

where

$$\xi_{it} = \tanh^{-1}\left(W_{it} - \mathrm{logit}^{-1}(\bar{\eta}_{it})\right), \tag{20}$$

denotes the innovation of $\bar{\eta}_{it}$ with respect to the observation W_{it}. The indicator A_{itl} terminates the memory of the process whenever the site has been closed between weeks $t-l$ and t. Since binary random variables give infinite values in the response space ($\mathrm{logit}(V_{it}) = \pm\infty$), we make comparisons in the input-space $[0,1]$ using the transformation $\mathrm{logit}^{-1}(\bar{\eta}_{it})$ first, and then transform differences back using $\tanh^{-1}(\cdot)$ which allows for negative values. Observe also that whenever $\mathrm{logit}^{-1}(\bar{\eta}_{it}) \approx W_{it}$, we have that $\xi_{it} \approx 0$ and we default to the best-guess estimate; meanwhile, larger discrepancies between $\mathrm{logit}^{-1}(\bar{\eta}_{it})$ and W_{it} give progressively larger values of ξ_{it} causing larger corrections. Through model estimation, we find an exponential decay in the auto-regressive coefficients; hence, we let

$$\phi_l = \phi^{\mathrm{scale}} \exp\left(\phi^{\mathrm{rate}}(l-1)\right), \tag{21}$$

and estimate only the two parameters ϕ^{scale} and ϕ^{rate}.

3.3 Joint forecasting

We now explain the steps and underlying assumptions used to derive joint forecasts using both models for lice abundance and treatment. Let $f(Y_{i,t+1}^\star \mid \mathcal{F}_t, V_{i,t+1})$ denote the conditional distribution for lice abundance in stage $\star$ at site i and week $t+1$ (as defined in Section 3.1) and let $f(W_{i,t+1} \mid \mathcal{F}_t)$ denote the conditional distribution of the occurrence of treatment (as defined in Section 3.2). For simplicity, we let $Y_{it} = (Y_{it}^{\mathrm{ST}}, Y_{it}^{\mathrm{MB}}, Y_{it}^{\mathrm{AF}})$ denote the concatenation of all stages of lice.

Assuming the counting of lice in each stage are conditionally independent of each other (conditional on $\mathcal{F}_t$ and $V_{i,t+1}$), we determine their joint distribution as

$$f(Y_{i,t+1} \mid \mathcal{F}_t, V_{i,t+1}) = \prod_{\star \in \{\mathrm{ST},\mathrm{MB},\mathrm{AF}\}} f(Y_{i,t+1}^\star \mid \mathcal{F}_t, V_{i,t+1}). \tag{22}$$

The need for treatment W_{it} is resolved by assuming farmers only use mechanical treatments since, lately, the majority of treatments are mechanical (see Figure 3). Thus, the conditional distribution for treatment is determined as

$$f(V_{i,t+1} \mid \mathcal{F}_t) = f(V_{i,t+1} \mid W_{i,t+1}) f(W_{i,t+1} \mid \mathcal{F}_t), \tag{23}$$

where the (deterministic) distribution $f(V_{i,t+1} \mid W_{i,t+1})$ translates $W_{i,t+1}$ to $V_{i,t+1}$. Lice abundance dependens on treatment within the same week, and their joint distribution is then determined as

$$f(Y_{i,t+1}, V_{i,t+1} \mid \mathcal{F}_t) = f(Y_{i,t+1} \mid \mathcal{F}_t, V_{i,t+1}) f(V_{i,t+1} \mid \mathcal{F}_t). \tag{24}$$

Consequently, the realisation of treatment $V_{i,t+1}$ must always be determined *before* lice abundance $Y_{i,t+1}$. It follows that the one-week-ahead forecast distribution at each site is a two component mixture consisting of two cases: with treatment or without treatment, where the probability of treatment determines their respective weights.

Up to now, we have only addressed the forecast distribution at an individual site. For simplicity, let $\boldsymbol{Y}_t = \{Y_{it}\}_{i=1,\ldots,N}$ and $\boldsymbol{V}_t = \{V_{it}\}_{i=1,\ldots,N}$ denote the lice abundances and treatments across all sites. We assume that *within a given week*, the realisation of treatment and lice abundance are conditionally independent across sites. Their joint distribution is then expressed as

$$f(\boldsymbol{Y}_{t+1}, \boldsymbol{V}_{t+1} \mid \mathcal{F}_t) = \prod_{i=1}^{N} f(Y_{i,t+1}, V_{i,t+1} \mid \mathcal{F}_t). \tag{25}$$

Recall that lice abundance and treatment at each site depend on the history of (almost) all other sites since these are connected by currents that transport lice between them. To forecast one additional week

at a single site, we must condition on the realisations of $\boldsymbol{Y}_{t+1}$ and $\boldsymbol{V}_{t+1}$ in the previous week, whose distribution is determined as

$$f(Y_{i,t+2}, V_{i,t+2} \mid \mathcal{F}_t) = f(Y_{i,t+2}, V_{i,t+2} \mid \mathcal{F}_t, \boldsymbol{Y}_{t+1}, \boldsymbol{V}_{t+1}) f(\boldsymbol{Y}_{t+1}, \boldsymbol{V}_{t+1} \mid \mathcal{F}_t). \tag{26}$$

While the one-week-ahead prediction has a simple analytical expression, we must resort to joint simulation across all sites to forecast beyond a single week. By repeating the above steps, we may also forecast multiple weeks ahead.

To summarise, the steps required to simulate the joint development of treatment and lice abundance across all sites are:

1. For each site i:
   (a) Predict treatment probability $p_{i,t+1}$ conditional on information $\mathcal{F}_t$
   (b) Sample a treatment realisation $W_{i,t+1}$ using probability $p_{i,t+1}$ and set $V^{\text{mec}}_{i,t+1} = W_{i,t+1}$
   (c) Predict the expected lice abundances $\mu^{\star}_{i,t+1}$ for each stage $\star$, conditional on $\mathcal{F}_t$ and the realisation of $V_{i,t+1}$
   (d) Sample the outcome of each lice count $n_{i,t+1}Y^{\star}_{i,t+1}$ given its expectation $n_{i,t+1}\mu^{\star}_{i,t+1}$
2. Let $\mathcal{F}_{t+1}$ include the sampled realisations of $\boldsymbol{Y}_{t+1}$ and $\boldsymbol{V}_{t+1}$ across all sites, as well as the past information set $\mathcal{F}_t$
3. Repeat the previous two steps for $t+2, t+3, \ldots, t+k$

By repeating the simulation scheme above, we can derive an empirical k-week-ahead forecast distribution $\hat{f}(\boldsymbol{Y}_{t+k}, \boldsymbol{V}_{t+k} \mid \mathcal{F}_t)$ conditional on all information $\mathcal{F}_t$ up to week t.

The reason for finding the joint forecast distribution is to infer different sites' exposure to future treatments. More importantly, we are also concerned about the *heterogeneity* in their risk exposure, which guides where to harvest within a portfolio of sites. As a proxy for risk exposure, we use the above simulation scheme to derive the empirical distribution of the aggregated future treatment count at each site within a k-week horizon:

$$R^k_{it} = \sum_{s=1}^{k} W_{i,t+s} \mid \mathcal{F}_t. \tag{27}$$

Meanwhile, risk management models apply the distribution $\hat{f}(\boldsymbol{Y}_{t+k}, \boldsymbol{V}_{t+k} \mid \mathcal{F}_t)$ directly. Information on future treatments can be incorporated into harvest plans by minimising overall exposure among multiple sites. Since harvest plans are made collectively, these must trade off which fish groups are to be exposed the longest, and such considerations may also be combined with information about the health state of the fish. This is investigated further in Section 4.3.

3.4 Estimation

The lice and treatment models were fitted separately by maximizing likelihood estimation (MLE). The log-likelihoods were implemented with the assistance of Template Model Builder (TMB), a free and open-source R package specifically designed for estimating complex non-linear models (Kristensen et al., 2016). The various parameter constraints are handled by maximizing a re-parameterised version of the log-likelihood, which varies over a set of unconstrained parameters. For instance, the constraint $\rho > 0$ for baseline weekly mortality is parameterized as $\rho = \exp(u)$, where u is an unconstrained parameter. TMB provides the gradient and Hessian computed using automatic differentiation (Fournier et al., 2012) at machine precision, and this is used in the R-routine *nlminb* for optimization. The C++ template functions have been made available online.[3] Standard deviations are computed using the Delta-method from the gradient and Hessian evaluated at the optimal parameter estimates.

Estimation is performed on data in the period 2018–2021 due to structural instabilities that stabilise by 2018. This is mainly attributed to the decline in the efficacy of medical treatments and changes in the counting dispersion of adult female lice. The transition was found by performing rolling estimation on fixed-width, two-year time windows. The Supplementary Material describes these observations in more detail.

[3] [Will be provided after publication]

4 Results and validation

We emphasize the forecasting ability of the joint model, which is analysed in Section 4.3. Additionally, we do in-sample analysis in Section 4.1 and assess the forecasting abilities of both the lice model and the treatment model in Section 4.2.

4.1 In-sample analysis

4.1.1 Parameter estimates

Estimates of all parameters for both models are given in Table 1 as well as their standard deviation derived by the Delta-method. For the lice model, we find that:

- The dispersion parameters $\nu^\star$ are incrementally increasing for $\star \in \{\text{ST}, \text{MB}, \text{AF}\}$, which means counting precision gets incrementally better for later stages. This corresponds well with reported experiences that adult female lice are easier to identify than those in the stationary stage (Thorvaldsen et al., 2019).
- The baseline weekly mortality rate is approximately $1 - \exp(\rho) = 5.47\%$ while the total effect of treatments are $1-\exp(\sum_k \delta_k^{\text{mec}}) = 38.8\%$ for mechanical and $1-\exp(\sum_k \delta_k^{\text{med}}) = 31.7\%$ for medical.
- We can infer that the memory from past innovations is sufficiently low at the last included lag due to the coefficient value of $\theta_{20} = \theta^{\text{scale}} \exp(\theta^{\text{rate}} \cdot 19) = 0.006$.

For the treatment model, we find that:

- Having a lower treatment limit (L_{it}) increases the odds of treatment by a multiplicative factor $\exp(\beta_L) = 2.16$, which is to be expected.
- The effect of lice on the probability of treatment is significant for both γ^{AF} and γ^{MB}. Meanwhile, the effect of stationary lice (Y_{it}^{ST}) on the probability of treatment was insignificant and removed during model selection.
- The memory from past innovations is sufficiently low since the last lagged auto-regressive coefficient has a value of $\phi_{20} = \phi^{\text{scale}} \exp(\phi^{\text{rate}} \cdot 19) = 0.037$.

Table 1: Parameter estimates for both models. Standard deviations of parameter estimates (SD) are derived by the Delta-method.

Lice model

Parameter	Estimate	SD
ν^{AF}	0.0880	0.0019
ν^{MB}	0.0535	0.0008
ν^{ST}	0.0253	0.0004
ι^{RU}	0.0446	0.0028
ι^{RS}	3.5767	0.1746
ι^{RN}	1.0337	0.0945
β^T	0.0584	0.0108
α	1.9760	0.0907
ρ	-0.0563	0.0064
δ_0^{mec}	-0.4911	0.0137
δ_0^{med}	-0.0000	0.0001
δ_1^{med}	-0.0935	0.0355
δ_2^{med}	-0.1187	0.0431
δ_3^{med}	-0.0575	0.0486
δ_4^{med}	-0.1113	0.0552
$\sum_k \delta_k^{\text{med}}$	-0.3810	0.0613
c	0.0065	0.0002
θ^{scale}	0.5709	0.0067
θ^{rate}	-0.2387	0.0060

Treatment model

Parameter	Estimate	SD
β_0	-6.5164	0.4789
β_L	0.7705	0.1190
γ^{AF}	3.9822	0.2264
γ^{MB}	3.0541	0.5056
$\gamma^{I,\text{AF}}$	-1.4766	0.8216
$\gamma^{I,\text{MB}}$	-2.2604	0.4635
π^{AF}	0.5774	0.0397
π^{MB}	0.1734	0.0292
ϕ^{scale}	0.4733	0.0459
ϕ^{rate}	-0.1336	0.0180

4.1.2 Sources of lice

Three sources contribute to lice recruits: neighbours (RN), within-site (RS), and unexplained sources (RU). When also accounting for development time and mortality rates, we can decompose the sources of lice by analysing each of their contributions to $\bar{\mu}^{\star}_{it} = \lambda^{\star,\mathrm{RS}}_{it} + \lambda^{\star,\mathrm{RN}}_{it} + \lambda^{\star,\mathrm{RU}}_{it}$ where

$$\lambda^{\star,\mathrm{RS}}_{it} = \kappa_{it0} \sum_{l=1}^{20} A_{itl} \kappa_{itl} r^{\star}_{itl} \hat{Y}^{\mathrm{RS}}_{i,t-l}, \tag{28a}$$

$$\lambda^{\star,\mathrm{RN}}_{it} = \kappa_{it0} \sum_{l=1}^{20} A_{itl} \kappa_{itl} r^{\star}_{itl} \hat{Y}^{\mathrm{RN}}_{i,t-l}, \tag{28b}$$

$$\lambda^{\star,\mathrm{RU}}_{it} = \kappa_{it0} \sum_{l=1}^{20} A_{itl} \kappa_{itl} r^{\star}_{itl} \hat{Y}^{\mathrm{RU}}_{i,t-l}. \tag{28c}$$

We do an in-sample analysis of this decomposition by averaging each source over all active sites in a given week to get a time-development of relative contributions: $\left(\sum_{i=1}^{N} A_{it}\lambda_{it}\right) / \left(\sum_{i=1}^{N} A_{it}\right)$. These are illustrated in Figure 10 in terms of average and normalised weekly quantities.

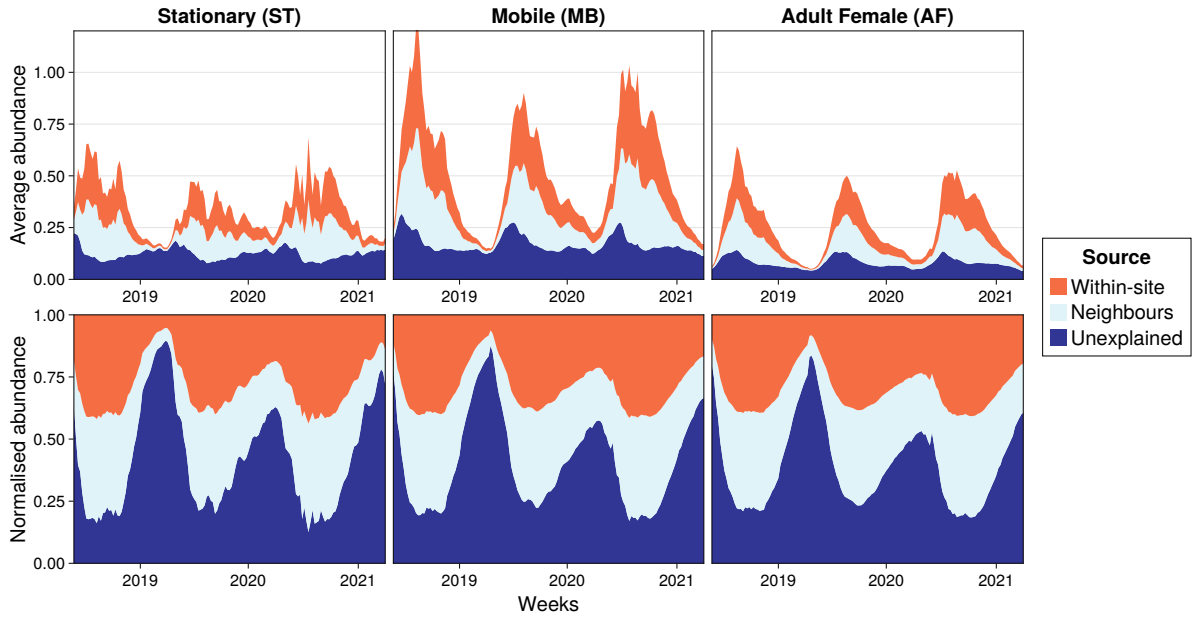

Figure 10: Sources that contribute to expected lice abundance at different stages, averaged over sites.

First, we see from Figure 10 a clear seasonal pattern due to temperature variations. Second, in late winter periods we observe that the explainable sources of lice reaches almost zero. This may be explained by increased time until infection and increased development time (due to low temperatures) such that the aggregated mortality rate (over longer time periods) is higher. Infectivity is also lower at low temperatures (see Figure 8). Third, we observe that there seems to be an almost equal contribution to lice abundance from neighbours as from within the same site. Lastly, we observe that the overall level of mobile lice is higher compared to the other stages; hence, the model reflects that this stage includes adult male lice in addition to pre-adults from both genders. This is also consistent with higher levels of mobile lice in the data.

4.1.3 Treatment patterns

While the auto-correlation function is a useful diagnostics tool for continuous ARMA processes, the *auto-persistence function* (APF) is a valuable counterpart for binary time-series data (Startz, 2008). The APF is stated as the conditional probability of future treatments:

$$\mathrm{APF}^0_k = P\left(V_{i,t+k} \mid V_{it} = 0\right), \qquad \mathrm{APF}^1_k = P\left(V_{i,t+k} \mid V_{it} = 1\right). \tag{29}$$

Figure 11 illustrates the empirical estimates of the APFs along with the corresponding in-sample estimates from the treatment model; both in terms of explanatory variables ($\bar{\eta}$) and using the correction by auto-regression (η).

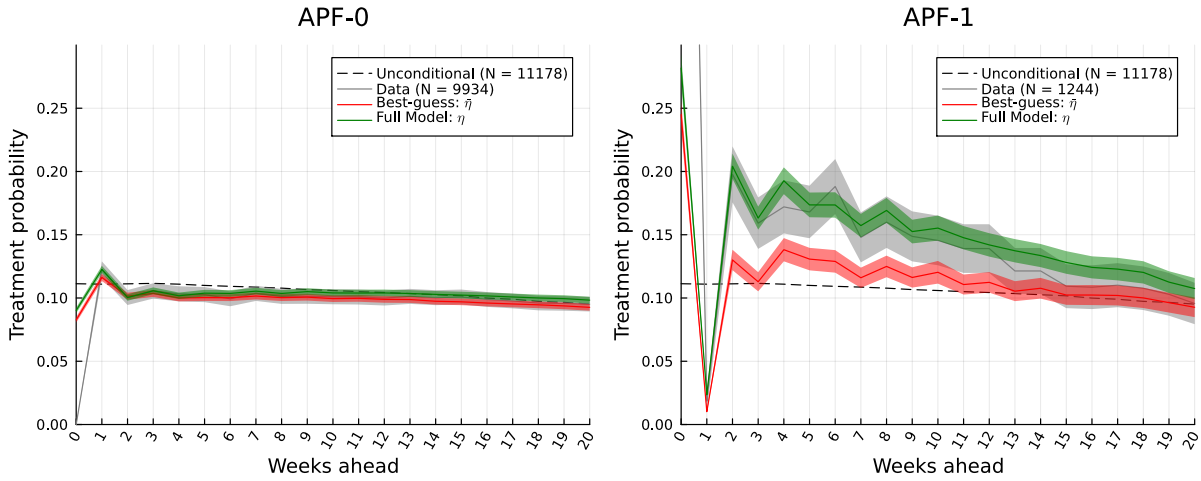

Figure 11: Auto-persistence function (APF) illustrating the empirical probability of future treatments conditional on treatment in the current week (APF-0 = No treatment, APF-1 = Treatment) in 20-week time windows of active weeks. Corresponding probabilities using in-sample model estimates with explanatory variables ($\bar{\eta}$) and with an added auto-regressive correction (η) are also shown. Ribbons give 95% confidence intervals.

From APF-1 in Figure 11, we see that treatments in the week following a week of treatment is very unlikely; however, this is also enforced in the data processing (see Section 2.1). APF-1 also shows a clear tendency that there is an increased probability of treatments from the second consecutive week and on-wards, given there is treatment in the current week. This implies that treatments tend to occur close in time, and that this effect decays over time. There are no notable characteristics in APF-0, neither in the data nor in the model predictions.

We may interpret the model-inferred value for logit($\bar{\eta}$) (which uses only explanatory variables) as the conditional probability of treatment, corrected for the effect of observed lice abundance. We observe from APF-1 in Figure 11 that this gives predictions that are slightly better than the baseline unconditional probability of treatment, but underestimates the conditional probability of treatment as observed in the data. Meanwhile, the full model (logit(η)) fits more closely to the empirical data and successfully corrects for the discrepancy between logit($\bar{\eta}$) and the data. This suggests lice abundance cannot predict treatments alone and that auto-regressive relations on previous treatments are able to explain additional effects.

On average, the treatment model assigns a 28.2% probability of treatment in weeks of treatment, while assigning a 9.0% probability of treatment in non-treatment weeks (from model estimates at week zero in Figure 11). A naïve model alternative would be an unconditional Bernoulli distribution specified by the average probability of treatment, which is 11%. Clearly, our model improves on this. Note that these numbers are derived within the sub-selection of weeks that start off all 20 week time windows where a site is active.

4.2 One-week-ahead forecasting ability

Risk management models use distributional forecasts as input directly; hence, we emphasize the forecast distributional fit of the current models. For this purpose, we use the probability integral transform (PIT) as suggested by Dawid (1984). The idea is to assess the percentile of the data within their forecast distribution: A correct diagnostic should resemble a uniform distribution of percentiles with the interpretation that the data resemble samples from their respective forecasting distributions. For discrete distributions, the PIT must be corrected to give a uniform distribution of percentiles (see Czado et al., 2009). By estimation up to time t, we find the forecast distribution for time $t+1$ at each site and evaluate the observed data within these (out-of-sample). The one-week-ahead forecast distributions are given analytically by (1) and (17), where we use the actual one-week-ahead treatment outcome to infer the distribution for lice abundance. This is repeated by rolling estimation, and evaluations are

combined across all sites (totally 67 weeks and 5576 active site-weeks). The resulting PITs are illustrated in Figure 12.

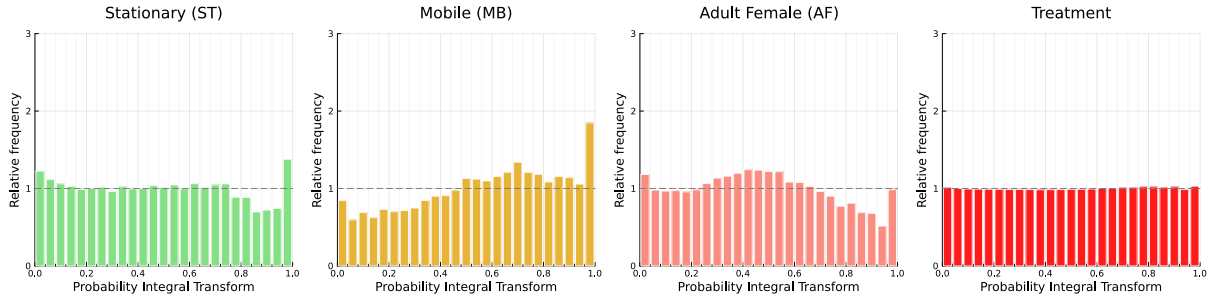

Figure 12: Corrected Probability Integral Transforms for forecast distributions for lice stages, and treatment probabilities at all sites.

The forecast distribution of treatments and stationary (ST) lice give PITs that are very close to uniformity. Still, the PITs for mobile (MB) and adult female (AF) lice indicate that the corresponding forecast distributions are slightly biased (by being non-centred) and over-dispersed (inverted U-shape), while some data points are more extreme than expected (inflated probabilities at the ends).

The irregularities in the forecast distributions of mobile (MB) and adult female (AF) lice have some possible explanations. First, the forecast distribution of adult female (AF) lice has a tendency to predict too large levels of lice. It is known that there are clear incentives for farmers to under-count adult female lice since this leads to costly treatments (Jeong et al., 2023), which might explain this bias. Furthermore, we also see a corresponding opposite bias in the predictions of mobile lice (MB); predicted lice levels are too low. These opposite biases may be related since there is no individual scaling between those quantities in the model, meaning the (incentive) bias in adult females may be the cause of both: predictions for MB are scaled down to account for the bias in AF. Various attempts to correct this bias typically lead to poor predictions of treatments in the joint model; hence, we do not correct this explicitly. Second, some data points are more extreme than expected, which is particularly pronounced for mobile lice (MB). This suggests the model does not capture some very rapid increases in lice abundance. We note there are occurrences of extreme lice abundance (above 5–10) in the data that are outside the range of all other counts, which may also explain some of these outliers. The GARMA correction also adjusts the overall level of forecasts in consecutive time steps after an outlier, and only an initial high lice abundance should cause large prediction errors.

We illustrate one-week-ahead out-of-sample predictions of the respective models in Figure 13 for an example site, along with actual data and prediction intervals of 50% and 95% confidence (illustrations for additional sites are provided in the Supplementary Material). First, we observe in Figure 13 that the forecast distributions for lice abundance seem to follow the data closely and replicate the dynamics well. For the treatment model, there is a close correspondence between higher predicted treatment probability and weeks of actual treatments; hence, it gives a meaningful distinction between weeks of lower and higher probability of treatment. In this particular case, the predicted probabilities in weeks of actual treatments are in the range of 20–40%, meaning there is still some ambiguity to the specific week a treatment will take place. Lastly, we observe that the forecast distributions for lice abundance show tendencies of over-dispersion since more data points are within the 50% prediction intervals than expected. A single site is not sufficient to draw the conclusion of over-dispersion but this is further supported by the PITs in Figure 12.

4.3 Long-term joint forecasting

The primary aim of the joint forecasting scheme is to predict the total number of treatments within a k-week-ahead horizon, denoted R_{it}^k in (27), since this is a major driver of costs and biological risk. The expected number of future treatments is an interesting measure of risk but variability is also important since it reflects whether meaningful deviations from the expected value may occur; production plans may need to account for such variability to imporve their performance.

First, we validate the forecasting ability of R_{it}^k. Figure 14 shows out-of-sample prediction intervals of aggregated future treatments within a 20-week horizon (starting April 2021) along with actual treatment counts during the same period. We see most data points are between these 95% prediction intervals

Figure 13: Out-of-sample one-week-ahead forecast distributions as prediction intervals and probability of treatment for an example site.

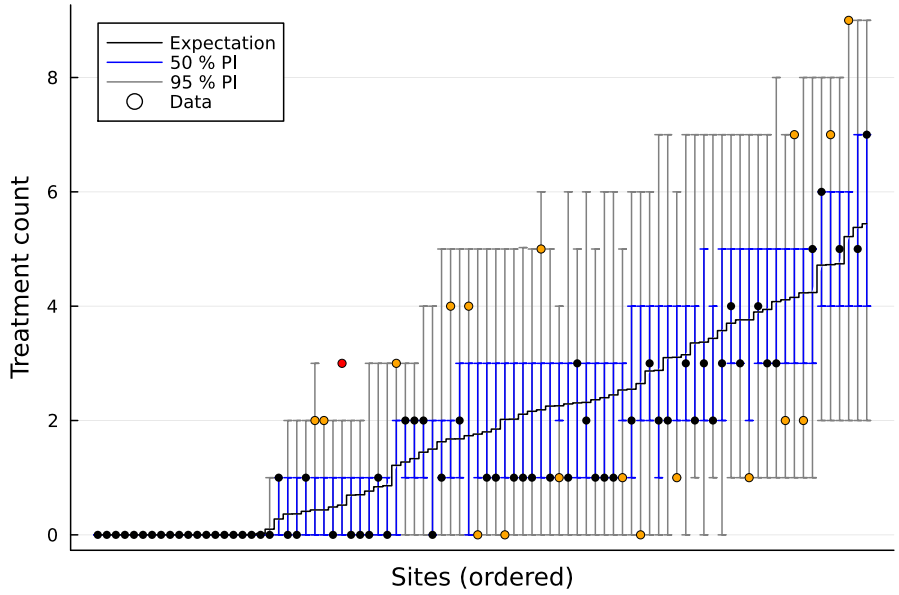

Figure 14: Forecasted 50% and 95% prediction interval for treatment count within a 20-week horizon. Sites on the first axis are sorted by average forecasted treatment counts and non-active sites are removed. The centers of the prediction intervals are averages. Observations that are outside the 50% prediction interval are colored orange, while those that are also outside the 95% prediction interval are colored red.

(specifically, 98.8% are within). We also notice that these prediction intervals are relatively wide, meaning there is large variability in the need for future treatments. Still, the joint forecasting model reflects this variability well.

4.3.1 Heterogeneity in risk exposure

For harvest planning and risk management, heterogeneity between sites is particularly interesting to guide the decision of *where* to harvest first. Farmers want to minimise overall risk exposure until a fish group is harvested subject to operational constraints that require different groups to be harvested in sequence. We do a simplified quantification of lice-induced losses to illustrate why heterogeneity in risk exposure matters. We account both for the direct cost of treatment and the indirect cost from loss in growth and increased fish mortality, which is compared to having no treatments at all. These costs are then scaled by the sales value of fish to get a relative measure of lice-induced loss. Additional details in these estimates are provided in the Supplementary Material. Figure 15 illustrates the distribution of site-specific losses within a 20-week horizon as kernel densities.

First, we observe from Figure 15 that the majority of sites have an expected loss between 0% and 21%, with an overall average exposure of 7.60% loss. Already, this constitutes very considerable amounts while the tail end of the worst loss distributions also reaches up to 40%. This quantification of risk can be very valuable since such estimates of still unrealised losses can be acted on in advance of their occurrence. The aim of the forecasting tool developed in this paper is to plan for such risk exposure proactively instead of only reacting to on-going developments. Second, we see that different sites have widely different levels of exposure that constitute very considerable amounts. For example, the probability of more-than 10% losses is significantly different among sites. Farmers may act on this information by prioritising to harvest the more exposed fish groups earlier than others.

5 Discussion and concluding remarks

Constructing a model for the given data has shown particularly challenging due to a relatively low signal-to-noise ratio, and an important consideration to ensure long-term forecasting ability has been to avoid over-fitting. The main strategy to prevent over-fitting has been to incorporate as much existing knowledge about lice dynamics as possible while also reducing the number of parameters. Model selection using in-sample criteria like likelihood, AIC and BIC has mostly been unhelpful, while building the model based on physical intuition has shown to be more effective. Still, the structure of the final model acknowledges

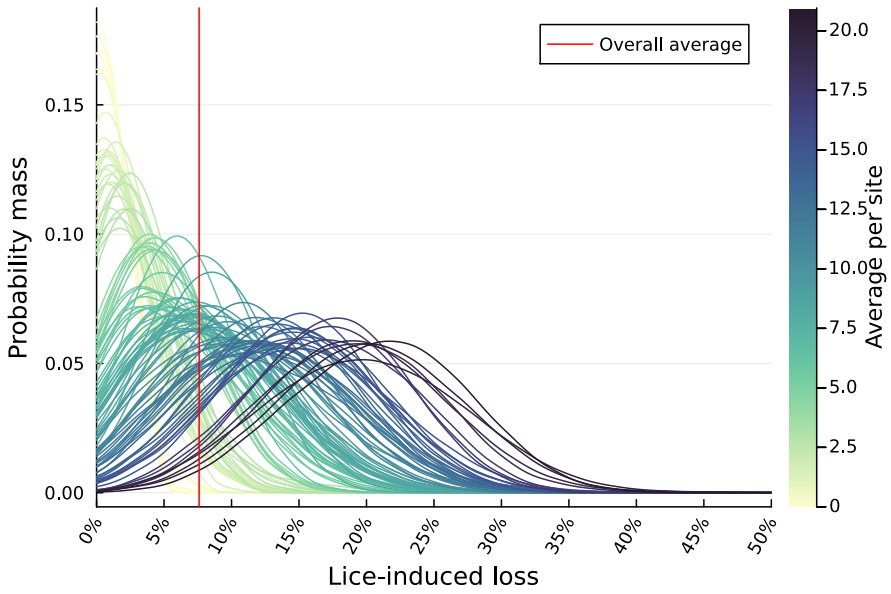

Figure 15: Site-specific risk exposure related to direct and indirect costs of lice treatments. Based on a 20-week forecast of treatment frequency on active sites. Illustrated as kernel densities for each site.

that not every aspect of lice dynamics can be explained by known causal mechanisms since it also requires an observation-based correction to more recent data (through GARMA models).

Incorporation of spatial effects is essential to determine heterogeneity in risk-exposure among sites since we already know this to be of major importance. In fact, the only source of heterogeneity in the joint model is stream patterns which, in turn, drive treatments at each site. Explicit use of stream patterns to derive heterogeneity between sites facilitates better trust that the model can performs well on future data and, due to its strong physical motivation, there is also reason to believe it performs well in new situations if sites are moved, removed or created. We emphasize trustworthiness since the model is to be applied to risk management, and validation of the joint model shows a good fit to the historical data of treatments that also replicates past heterogeneity between sites.

We model treatments as a stochastic process because the ambiguity is too high for when treatment limits are truly surpassed, likely due to the high counting dispersion. Simply determining treatments based on whether simulated lice counts are above the limit does not work very well in the forecasting model. In Figure 11, we show that even a regression model on several explanatory variables that includes lice abundance is insufficient to explain treatment patterns in the data. Only once auto-regressions on past treatments are added do we replicate what is observed in the data. An interpretation of why auto-regressive relations matter is that farmers attain more information than is contained in the lice counts. For example, farmers may observe the fish visually without taking them out for counting, or they may have experience with how fast lice abundance develops at their site. Such cues are not reported and cannot be used in a model; however, auto-regressive relations on past treatments can capture similar relations to give an effective prediction model. In other words, by modelling auto-regressive relations we assume more information goes into treatment decisions than is contained in the lice counts themselves.

Industry know-how suggests that the effectiveness of treatments is in the range of 90%; thus, our parameter estimates (approximately 30–40%) show less effectiveness than expected. Aldrin et al. (2019) also found this effect to be in a similar range to ours. There are two possible explanations to this large discrepancy. First, it may be due to a spurious correlation between treatments and lice abundance. Treatments co-occur with increased lice abundance but treatments are *not* the cause of more lice; rather, the opposite is true. Since it is difficult to separate the causal effect of treatments on lice abundance from their co-occurrence, the prediction power for lice abundance may counter-intuitively be improved by a lower effect of treatment. For the same reason, we constrain the sign on the effectiveness of treatment to avoid getting negative values. Second, one may speculate whether lice can end up in the water before (or during) treatments and then return to the host afterwards. Even if the treatment itself has a 90% efficacy, the overall effect must be lower if the louse later returns to its host. Our model estimates would only reflect the overall effect of treatment, which should then also be lower than the expected 90% if lice later return to their hosts.

In future research, it may be advantageous to incorporate higher resolution stream patterns to account for known seasonal variations. This would require running longer hydrodynamic simulations, which is outside the scope of the current paper. Furthermore, future research may also validate the models using more precise machine vision counting which is currently being deployed in the industry. In principle, these could be applied directly without major alterations.

Acknowledgements

We thank Bjørn Ådlandsvik at the Norwegian Institute of Marine Research for providing hydrodynamic simulation data. The Norwegian Research Council supports both authors: The first author through the project "Optimal Risk-based Short-term Decision Making for Aquaculture" (#310001), and the second author through the Centre for Research-Based Innovation "Climate Futures" (#309562).

References

Abolofia, J., Asche, F. & Wilen, J. E. (2017). The Cost of Lice: Quantifying the Impacts of Parasitic Sea Lice on Farmed Salmon. *Marine Resource Economics*, *32*(3), 329–349. https://doi.org/10.1086/691981

Albretsen, J., Sperrevik, A. K., Staalstrøm, A., Sandvik, A. D., Vikebø, F. & Asplin, L. (2011). NorKyst-800 Rapport nr. 1 : Brukermanual og tekniske beskrivelser. *48 s.* Retrieved June 22, 2021, from https://imr.brage.unit.no/imr-xmlui/handle/11250/113865
Accepted: 2011-07-13T11:53:51Z

Aldrin, M., Huseby, R., Stien, A., Grøntvedt, R., Viljugrein, H. & Jansen, P. (2017). A stage-structured Bayesian hierarchical model for salmon lice populations at individual salmon farms – Estimated from multiple farm data sets. *Ecological Modelling*, *359*, 333–348. https://doi.org/10.1016/j.ecolmodel.2017.05.019

Aldrin, M., Jansen, P. A. & Stryhn, H. (2019). A partly stage-structured model for the abundance of salmon lice in salmonid farms. *Epidemics*, *26*, 9–22. https://doi.org/10.1016/j.epidem.2018.08.001

Aldrin, M., Storvik, B., Kristoffersen, A. B. & Jansen, P. A. (2013). Space-Time Modelling of the Spread of Salmon Lice between and within Norwegian Marine Salmon Farms (M. Krkosek, Ed.). *PLoS ONE*, *8*(5), e64039. https://doi.org/10.1371/journal.pone.0064039

Asplin, L., Albretsen, J., Johnsen, I. A. & Sandvik, A. D. (2020). The hydrodynamic foundation for salmon lice dispersion modeling along the Norwegian coast. *Ocean Dynamics*, *70*(8), 1151–1167. https://doi.org/10.1007/s10236-020-01378-0

Bang Jensen, B., Qviller, L. & Toft, N. (2020). Spatio-temporal variations in mortality during the seawater production phase of Atlantic salmon (Salmo salar) in Norway. *Journal of Fish Diseases*, *43*(4), 445–457. https://doi.org/10.1111/jfd.13142

Benjamin, M. A., Rigby, R. A. & Stasinopoulos, D. M. (2003). Generalized Autoregressive Moving Average Models. *Journal of the American Statistical Association*, *98*(461), 214–223. https://doi.org/10.1198/016214503388619238

Czado, C., Gneiting, T. & Held, L. (2009). Predictive Model Assessment for Count Data. *Biometrics*, *65*(4), 1254–1261. https://doi.org/10.1111/j.1541-0420.2009.01191.x

Dawid, A. P. (1984). Present Position and Potential Developments: Some Personal Views: Statistical Theory: The Prequential Approach. *Journal of the Royal Statistical Society. Series A (General)*, *147*(2), 278. https://doi.org/10.2307/2981683

Elghafghuf, A., Vanderstichel, R., Hammell, L. & Stryhn, H. (2020). Estimating sea lice infestation pressure on salmon farms: Comparing different methods using multivariate state-space models. *Epidemics*, *31*, 100394. https://doi.org/10.1016/j.epidem.2020.100394

Elghafghuf, A., Vanderstichel, R., St-Hilaire, S. & Stryhn, H. (2018). Using state-space models to predict the abundance of juvenile and adult sea lice on Atlantic salmon. *Epidemics*, *24*, 76–87. https://doi.org/10.1016/j.epidem.2018.04.002

Fournier, D. A., Skaug, H. J., Ancheta, J., Ianelli, J., Magnusson, A., Maunder, M. N., Nielsen, A. & Sibert, J. (2012). Ad model builder: Using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. *Optimization Methods and Software*, *27*(2), 233–249. https://doi.org/10.1080/10556788.2011.597854

Hamre, L. A., Eichner, C., Caipang, C. M. A., Dalvin, S. T., Bron, J. E., Nilsen, F., Boxshall, G. & Skern-Mauritzen, R. (2013). The Salmon Louse Lepeophtheirus salmonis (Copepoda: Caligidae) Life Cycle Has Only Two Chalimus Stages. *PLOS ONE*, *8*(9), e73539. https://doi.org/10.1371/journal.pone.0073539

Hamre, L. A., Bui, S., Oppedal, F., Skern-Mauritzen, R. & Dalvin, S. (2019). Development of the salmon louse Lepeophtheirus salmonis parasitic stages in temperatures ranging from 3 to 24°C. *Aquaculture Environment Interactions*, *11*, 429–443. https://doi.org/10.3354/aei00320

Hellemo, L., Barton, P. I. & Tomasgard, A. (2018). Decision-dependent probabilities in stochastic programs with recourse. *Computational Management Science*, *15*(3-4), 369–395. https://doi.org/10.1007/s10287-018-0330-0

Iversen, A., Hermansen, Ø., Nystøyl, R. & Hess, E. J. (2017). *Kostnadsutvikling i lakseoppdrett – med fokus på fôr- og lusekostnader*. Nofima AS. Retrieved December 5, 2020, from https://nofima.brage.unit.no/nofima-xmlui/handle/11250/2481501
Accepted: 2018-02-01T14:47:27Z

Jansen, P. A., Kristoffersen, A. B., Viljugrein, H., Jimenez, D., Aldrin, M. & Stien, A. (2012). Sea lice as a density-dependent constraint to salmonid farming. *Proceedings of the Royal Society B: Biological Sciences*, *279*(1737), 2330–2338. https://doi.org/10.1098/rspb.2012.0084

Jensen, E. M., Horsberg, T. E., Sevatdal, S. & Helgesen, K. O. (2020). Trends in de-lousing of Norwegian farmed salmon from 2000–2019—Consumption of medicines, salmon louse resistance and non-medicinal control methods. *PLOS ONE*, *15*(10), e0240894. https://doi.org/10.1371/journal.pone.0240894

Jeong, J., Arriagada, G. & Revie, C. W. (2023). Targets and measures: Challenges associated with reporting low sea lice levels on Atlantic salmon farms. *Aquaculture*, *563*, 738865. https://doi.org/10.1016/j.aquaculture.2022.738865

Jeong, J. & Revie, C. W. (2020). Appropriate sampling strategies to estimate sea lice prevalence on salmon farms with low infestation levels. *Aquaculture*, *518*, 734858. https://doi.org/10.1016/j.aquaculture.2019.734858

Jonsbråten, T. W., Wets, R. J.-B. & Woodruff, D. L. (1998). A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, *82*(0), 83–106. https://doi.org/10.1023/A:1018943626786

King, A. J. & Wallace, S. W. (2012). *Modeling with Stochastic Programming*. Springer New York, NY. https://doi.org/10.1007/978-0-387-87817-1

Kristensen, K., Nielsen, A., Berg, C., Skaug, H. & Bell, B. (2016). TMB: Automatic differentiation and laplace approximation. *Journal of Statistical Software, Articles*, *70*(5), 1–21. https://doi.org/10.18637/jss.v070.i05

Kristoffersen, A. B., Qviller, L., Helgesen, K. O., Vollset, K. W., Viljugrein, H. & Jansen, P. A. (2018). Quantitative risk assessment of salmon louse-induced mortality of seaward-migrating post-smolt Atlantic salmon. *Epidemics*, *23*, 19–33. https://doi.org/10.1016/j.epidem.2017.11.001

Lovdata. (2008). Forskrift om drift av akvakulturanlegg (akvakulturdriftsforskriften) - Lovdata. Retrieved May 5, 2023, from https://lovdata.no/dokument/SF/forskrift/2008-06-17-822

Lovdata. (2012). Forskrift om bekjempelse av lakselus i akvakulturanlegg - Lovdata. Retrieved September 8, 2021, from https://lovdata.no/dokument/SF/forskrift/2012-12-05-1140

Lovdata. (2017). Forskrift om endring i forskrift om bekjempelse av lakselus i akvakulturanlegg - Lovdata. Retrieved July 5, 2023, from https://lovdata.no/dokument/LTI/forskrift/2017-03-06-275

Lovdata. (2018). Forskrift om endring i forskrifter om akvakultur for deres tilpasning til transport, oppbevaring, bruk og produksjon av rensefisk - Lovdata. Retrieved July 5, 2023, from https://lovdata.no/dokument/LTI/forskrift/2018-04-19-674

Mowi ASA. (2023). Integrated Annual Report 2022. https://mowi.com/wp-content/uploads/2023/03/Mowi-Integrated-Annual-Report-2022.pdf

Myksvoll, M. S., Sandvik, A. D., Albretsen, J., Asplin, L., Johnsen, I. A., Karlsen, Ø., Kristensen, N. M., Melsom, A., Skardhamar, J. & Ådlandsvik, B. (2018). Evaluation of a national operational salmon lice monitoring system—From physics to fish. *PLOS ONE*, *13*(7), e0201338. https://doi.org/10.1371/journal.pone.0201338

Samsing, F., Johnsen, I., Stien, L. H., Oppedal, F., Albretsen, J., Asplin, L. & Dempster, T. (2016). Predicting the effectiveness of depth-based technologies to prevent salmon lice infection using a dispersal model. *Preventive Veterinary Medicine*, *129*, 48–57. https://doi.org/10.1016/j.prevetmed.2016.05.010

Shapiro, A., Dentcheva, D. & Ruszczyński, A. P. (2014). *Lectures on stochastic programming: Modeling and theory* (Second edition). Society for Industrial and Applied Mathematics : Mathematical Optimization Society. https://doi.org/10.1137/1.9781611973433

Skern-Mauritzen, R., Sissener, N. H., Sandvik, A. D., Meier, S., Sævik, P. N., Skogen, M. D., Vågseth, T., Dalvin, S., Skern-Mauritzen, M. & Bui, S. (2020). Parasite development affect dispersal dynamics; infectivity, activity and energetic status in cohorts of salmon louse copepodids. *Journal of Experimental Marine Biology and Ecology*, *530–531*, 151429. https://doi.org/10.1016/j.jembe.2020.151429

Solberg, I., Finstad, B., Berntsen, H. H., Diserud, O. H., Frank, K., Helgesen, K. O., Jeong, J., Kristoffersen, A. B., Nytrø, A. V., Revie, C. W., Sivertsgård, R., Solvang, T., Sunde, L. M., Thorvaldsen, T., Uglem, I. & Mo, T. A. (2018). *Kartlegging og testing av metodikk for telling av lakselus og beregning av luseforekomst* (tech. rep.). Norsk Institutt for Naturforskning (NINA). Retrieved July 20, 2022, from https://brage.nina.no/nina-xmlui/handle/11250/2560950
Accepted: 2018-09-05T11:33:46Z

Startz, R. (2008). Binomial Autoregressive Moving Average Models With an Application to U.S. Recessions. *Journal of Business & Economic Statistics*, *26*(1), 1–8. https://doi.org/10.1198/073500107000000151

Stien, A., Bjørn, P. A., Heuch, P. A. & Elston, D. A. (2005). Population dynamics of salmon lice Lepeophtheirus salmonis on Atlantic salmon and sea trout. *Marine Ecology Progress Series*, *290*, 263–275. https://doi.org/10.3354/meps290263

Thorvaldsen, T., Frank, K. & Sunde, L. M. (2019). Practices to obtain lice counts at Norwegian salmon farms: Status and possible implications for representativity. *Aquaculture Environment Interactions*, *11*, 393–404. https://doi.org/10.3354/aei00323

Walde, C. S., Jensen, B. B., Pettersen, J. M. & Stormoen, M. (2021). Estimating cage-level mortality distributions following different delousing treatments of Atlantic salmon (*salmo salar*) in Norway. *Journal of Fish Diseases*, *44*(7), 899–912. https://doi.org/10.1111/jfd.13348

Walde, C. S., Stormoen, M., Pettersen, J. M., Persson, D., Røsæg, M. V. & Bang Jensen, B. (2022). How delousing affects the short-term growth of Atlantic salmon (Salmo salar). *Aquaculture*, *561*, 738720. https://doi.org/10.1016/j.aquaculture.2022.738720

Supplementary Material

1 Biological relations

We use several relations based on the biology of lice that has been derived from empirical laboratory experiments. Due to wider temperature ranges and a requirement to extrapolate outside the temperature ranges in the experiments, we have re-estimated expressions for these relations. The data has been extracted from plots in the original papers using graphing software.

1.1 Development time between stages and hatch rate

Lice development time has been examined experimentally by Hamre et al. (2019) in the temperature range $6°C$ to $21°C$, and development in the larval (LR) phase by Samsing et al. (2016) in the temperature range $5°C$ to $20°C$. While they report low or almost no development at $3°C$, there are also exist accounts of lice developing at temperatures down to $2°C$. For forecasting lice, we have temperatures in the range $-0.5°C$ to $24°C$, and require these expressions to extrapolate beyond the given temperatures, which is not possible using the expressions estimated by Hamre et al. (2019) and Samsing et al. (2016) since they give either negative or undefined development times. The data in (Hamre et al., 2019; Samsing et al., 2016) both exhibit second-order tendencies in development times and are required to be positive; hence, we choose an expression for development times in the form

$$\exp(a + bT + cT^2). \tag{1}$$

The newly fitted expressions coincide reasonably close to the original estimated expressions within their relevant temperature ranges. We use the same procedure to infer temperature dependent hatch rates from data in (Samsing et al., 2016).

1.2 Time window and infectivity

The time window of infectivity is defined by the time it takes for larvae to develop and be able to attach to a new host up until the time it runs out of nutrition and dies. We use data from Samsing et al. (2016) to estimate these functions using quantile regression at the 80th percentile to get a conservative estimate of the width of the time window. The time to develop into the larval stage at the 80th percentile is estimated to

$$D_L^{80}(T) = \exp\left(1.492041 - 0.027941T - 0.001701T^2\right), \tag{2}$$

and the width of the infestation window at the 80th percentile is estimated to

$$W^{80}(T) = \exp\left(0.450114 + 0.090447T - 0.005636T^2\right). \tag{3}$$

The lower and upper limits of the time window is then expressed as

$$W_l(T) = D_L^{80}(T), \tag{4}$$

$$W_u(T) = D_L^{80}(T) + W^{80}(T), \tag{5}$$

respectively.

Skern-Mauritzen et al. (2020) shown there are important interactions between age and temperature to determine infestation success. The temperature effect was examined experimentally by Samsing et al. (2016) in the temperature range $5°C$ to $20°C$ while the interaction between age and temperature was examined by Skern-Mauritzen et al. (2020) in the temperature range $5°$ to $15°C$. The expression for infectivity estimated by Skern-Mauritzen et al. (2020) seems to be over-fitted since it deviates from the

data by (Samsing et al., 2016) at 20°C and they also report concerns about over-fitting themselves. Hence, we estimate a new expression for infectivity which allows for extrapolation.

We use the time window $[W_l(T), W_u(T)]$ as a starting point, and specify a uni-modal function of age A within this range that has temperature dependent parameters. Then, infectivity takes the expression

$$I(A,T) = \begin{cases} C(T)\frac{[A]^{a(T)}(1-[A])^{b(T)}}{\mathcal{B}(a(T),b(T))} & \text{if } W_l(T) \leq A \leq W_u(T) \\ 0 & \text{otherwise,} \end{cases} \tag{6}$$

where A is age in weeks, $[A] = \frac{A-W_l(T)}{W_u(T)-W_l(T)}$ is the fractional age within the time window, and $\mathcal{B}(\cdot,\cdot)$ is the Beta function. This expression has been validated to adhere to the data from (Samsing et al., 2016) and has reasonable properties for extrapolation outside the given temperature range of the data. The temperature dependent parameters are estimated to

$$C(T) = \exp\left(-3.068236 + 0.447180T - 0.012281T^2\right), \tag{7}$$

$$a(T) = \exp\left(4.465295 - 0.619186T + 0.018293T^2\right), \tag{8}$$

$$b(T) = \exp\left(0.766574 + 0.133997T - 0.010022T^2\right), \tag{9}$$

and the function $I(A,T)$ is illustrated in the paper. For evaluation of infestation success for ages between $s-1$ and s weeks, we use the average infectivity

$$\bar{I}_s(T) = \int_{s-1}^{s} I(A,T)dA, \tag{10}$$

during the week that leads up to week s.

2 Structural instability

There is evidence of structural instability in the lice model for the given time-series data. Namely, there is a systematic drift in parameter estimates over time, implying changes in the data generation mechanism not captured by our model. To detect structural instability, we perform rolling estimation within two year time windows, and compare their respective parameter estimates. Rolling parameter estimates are illustrated in Figure 1 for the lice model, and in Figure 2 for the treatment model. Note that we fix the effect of cleaner fish for time windows that contain data after 2018 since there is no new data on cleaner fish after that point.

Foremost, we see systematic changes in the lice model generally occurring in the period 2016–2019, and some particular violations of structural stability include:

- There seems to be a decrease in dispersion of counts for adult female (AF) lice (by higher values of ν^{AF}). This may, among other things, be explained by more standardized counting practices (Solberg et al., 2018).
- We see an overall decrease in the effect of medical treatments up to the beginning of 2017. This is a known phenomenon caused by resistance towards the used substances (Jensen et al., 2020).
- There is an increasing effect of mechanical treatments up to 2017. This may be explained by accounts that better treatment practices learned in the industry during 2016–2018.

There are several regulatory changes during this period (Lovdata, 2016, 2017, 2018) that may explain some of these effects. In previous work, Aldrin et al. (2019) concluded that the treatment effect should be on the 1st lag, using data up to the end of 2016. Our analysis shows a shift occurred later and we, in contrast, also use a 0-lagged effect of treatment. It is mainly the introduction of mechanical treatments after 2017 that is the cause of this shift. If we do not distinguishing medical and mechanical treatment, we get severe structural instabilities in the timing of treatment effects.

For the treatment model, the main structural instability is in the effect of adult female lice through the parameters γ^{AF} and π^{AF} that seem to be increasing. This may be linked to the dispersion in these lice counts that has become lower during the same time period, meaning the counts of adult female lice provide better signals to infer treatments. There also seems to be an increasing effect of using past treatments to infer the probability of treatment, due to larger estimates of the parameter ϕ^{AF} over time.

We observe that most instabilities in parameter estimates have stabilised around 2018, which is why we present parameter estimates using data after 2018 in the paper.

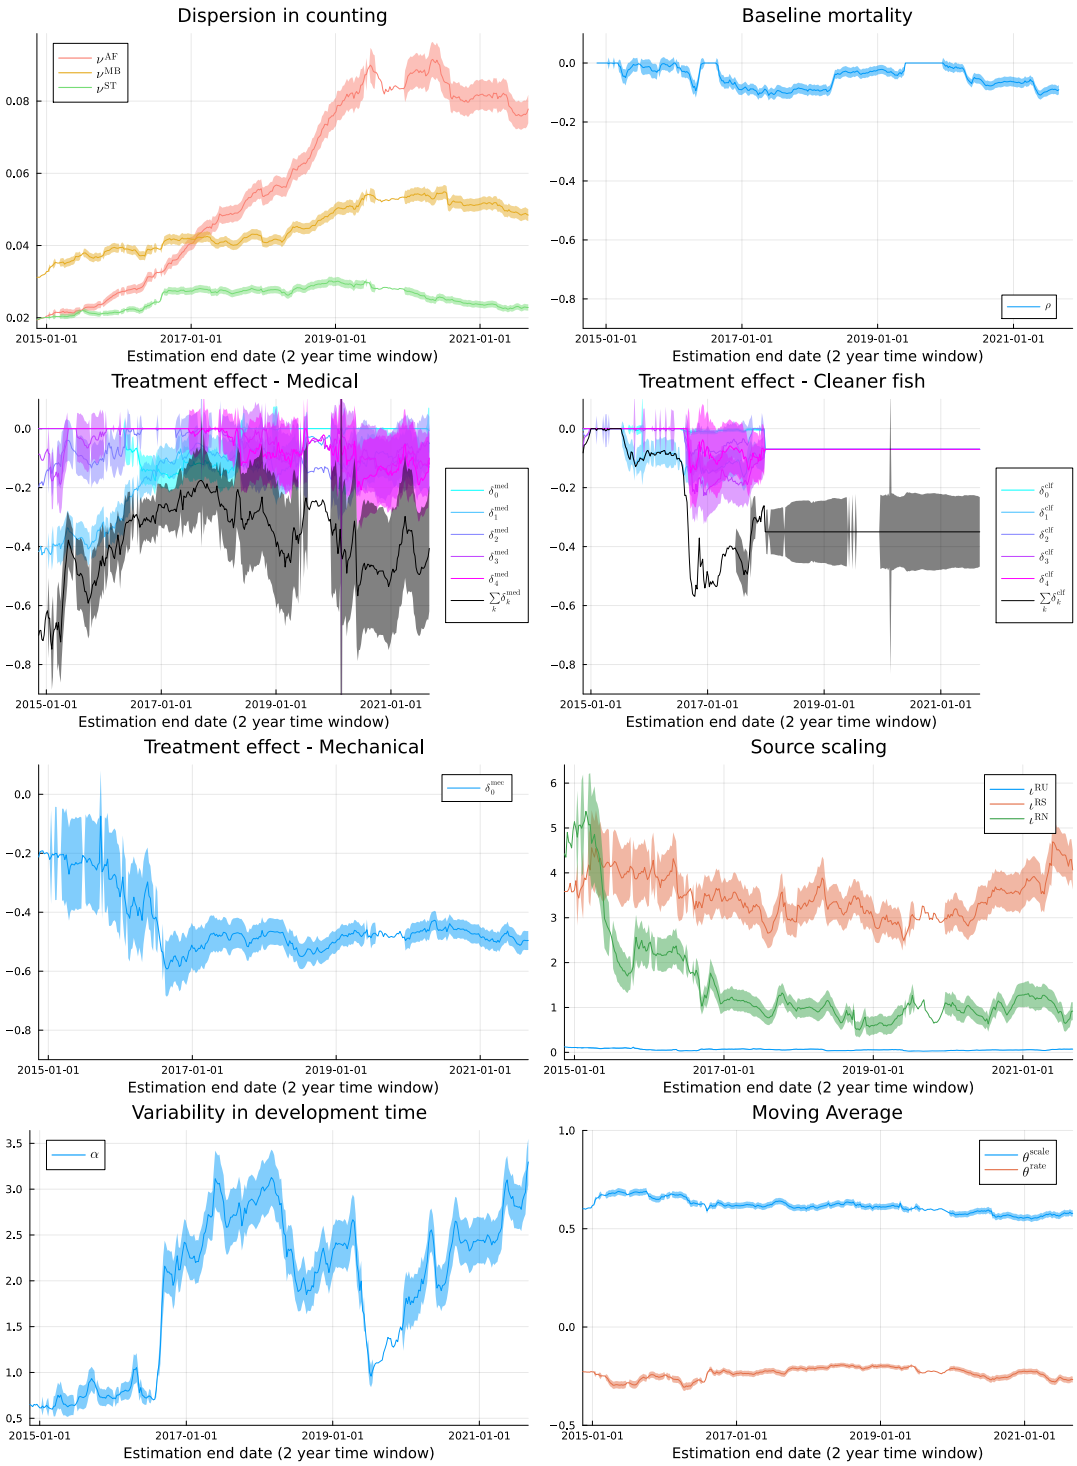

Figure 1: Lice model: Rolling parameter estimates within three year time windows. Ribbons give 95% confidence intervals derived by the Delta method.

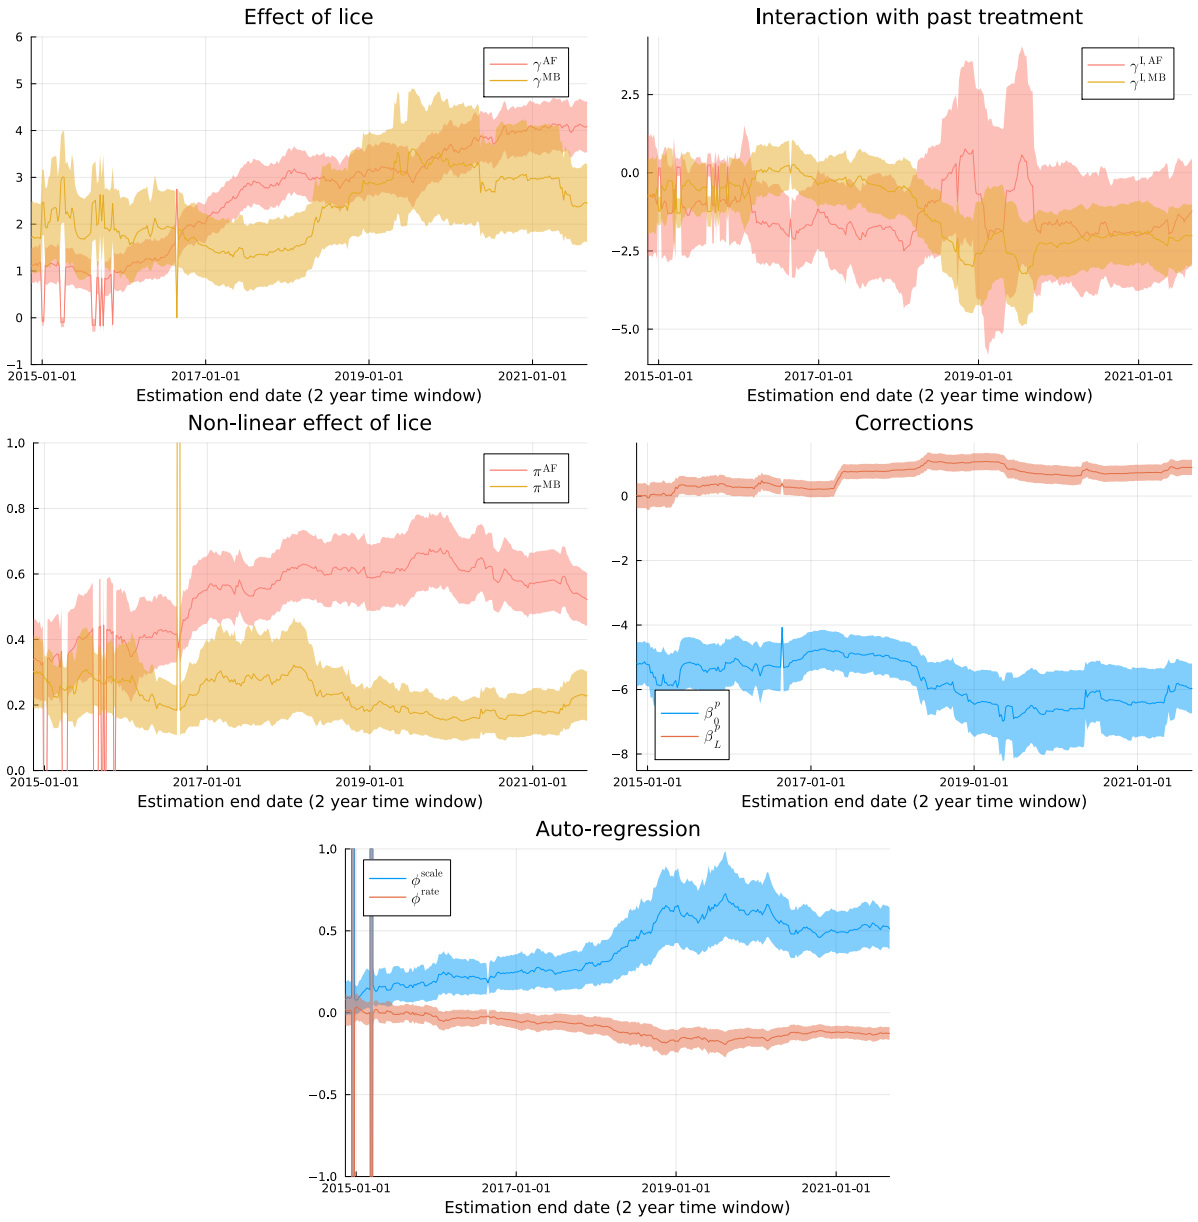

Figure 2: Treatment model: Rolling parameter estimates within three year time windows. Ribbons give 95% confidence intervals derived by the Delta method.

3 Quantifying the cost of lice treatments

We make some rough assumptions to quantify the cost of treatment. First, we assume we have a cage of fish at 2 kg and that we have a baseline daily growth rate of 0.5% given a temperature of $10°C$ (Skretting, 2012). Before a treatment, the fish are starved for 3–5 days to reduce stress and, additionally, the treatment itself cause reduced growth afterwards. Combined, these cause a approximate 32% decreased growth for 1–2 weeks (Walde et al., 2022). Weekly mortality rate increase approximately 1.17 percentage points as a result of mechanical treatments (Walde et al., 2021). The treatment operation in itself costs approximately 250 000 NOK per cage, which we distribute evenly across 200 000 fish. We assume an approximate sales price of 65 NOK/kg. To estimate the cost of lice treatments, we forecast growth and mortality for the next 20 weeks within two scenarios: (i) including the effect of treatments, or (ii) ignoring the effect of treatments. The cost of lice is estimated as the difference in revenue between these two scenarios. The relative measure of lice-induced loss equals the cost of lice divided by the revenue that includes treatments. More precise accounts of the operational cost of treatment are provided by Abolofia et al. (2017) and Iversen et al. (2017).

4 Forecasts distributions for additional sites

One-week-ahead forecast distributions for three additional sites are shown in Figure 3, 4 and 5.

Figure 3: Out-of-sample one week ahead forecast distributions as prediction intervals and probability of treatment.

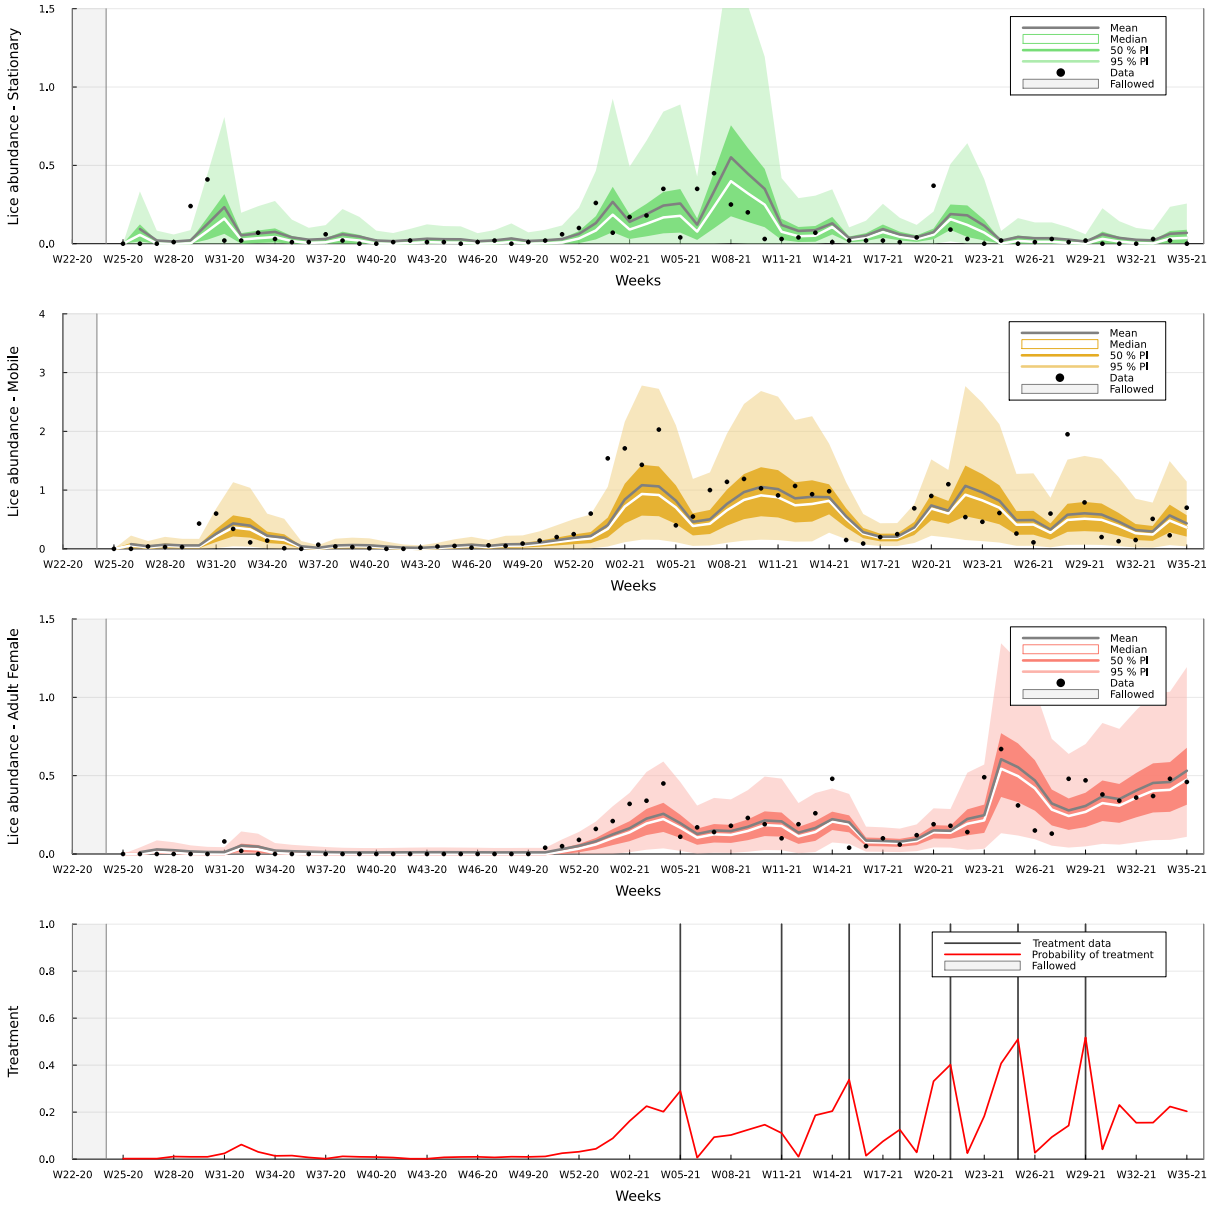

Figure 4: Out-of-sample one week ahead forecast distributions as prediction intervals and probability of treatment.

Figure 5: Out-of-sample one week ahead forecast distributions as prediction intervals and probability of treatment.

NORGES HANDELSHØYSKOLE
Norwegian School of Economics

Helleveien 30
NO-5045 Bergen
Norway

T +47 55 95 90 00
E nhh.postmottak@nhh.no
W www.nhh.no