Spatiotemporal modelling of marine movement data using Template Model Builder

Marie Auger-Méthé*, Christoffer M. Albertsen, Ian D. Jonsen, Andrew E. Derocher, Damian C. Lidgard, Katharine R. Studholme, W. Don Bowen, Glenn T. Crossin, Joanna Mills Flemming

*Corresponding author: auger-methe@dal.ca

Supplements to Marine Ecology Progress Series 565:237-249

S1 Additional information for case study 1

As discussed in the main text, we used `TMB` and `bsam` to fit the DCRW to the Argos tracks of four polar bears. This appendix contains additional information on the data used, the model applied, and the results of both the polar bear analysis and associated simulation study.

S1.1 Extra polar bear data information

Most Argos locations from the bears were from the lower quality categories (A-B). The number of locations in each category is presented in Table S1.1.

Table S1.1: Number of locations in each Argos category for the four polar bear movement tracks.

ID	3	2	1	0	A	B	Start date	End date
PB1	14	36	36	32	676	1125	2009-04-20	2010-04-21
PB2	19	41	72	64	1292	2152	2009-04-19	2011-06-10
PB3	8	12	34	43	597	1040	2011-04-22	2012-06-26
PB4	10	17	39	34	580	924	2010-04-25	2011-06-29

S1.2 Priors used in `bsam`

`bsam` uses a Bayesian framework to fit the DCRW to data. As such, the full DCRW model in `bsam` includes priors (described in Table S1.2). We used the default priors from `bsam`. To explore how these priors may influence the analysis, we used `rjags` to sample 5000 values from each prior. The histogram of these samples are presented in Fig. S1.1. From both Table S1.2 and Fig. S1.1 we can see that θ and γ have an equal probability throughout their parameter space. In contrast, ψ and ρ have much higher probability of values very close to 0, while $\sigma_{\epsilon,lon}$ and $\sigma_{\epsilon,lat}$ have peaks around 0.8. Some of these priors may have helped the estimation, for example the peak in probability of ρ close to 0 may help the overall estimation of the process covariance matrix $\mathbf{\Sigma}_\epsilon$.

Table S1.2: Priors used in `bsam`.

Parameter	Prior distribution
γ	$\text{Beta}(1,1)$
θ	$\text{U}(-\pi,\pi)$
ψ	$\exp(\text{U}(-10,10))$
Σ_ϵ	$\text{W}^{-1}(\left[\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right],2)$

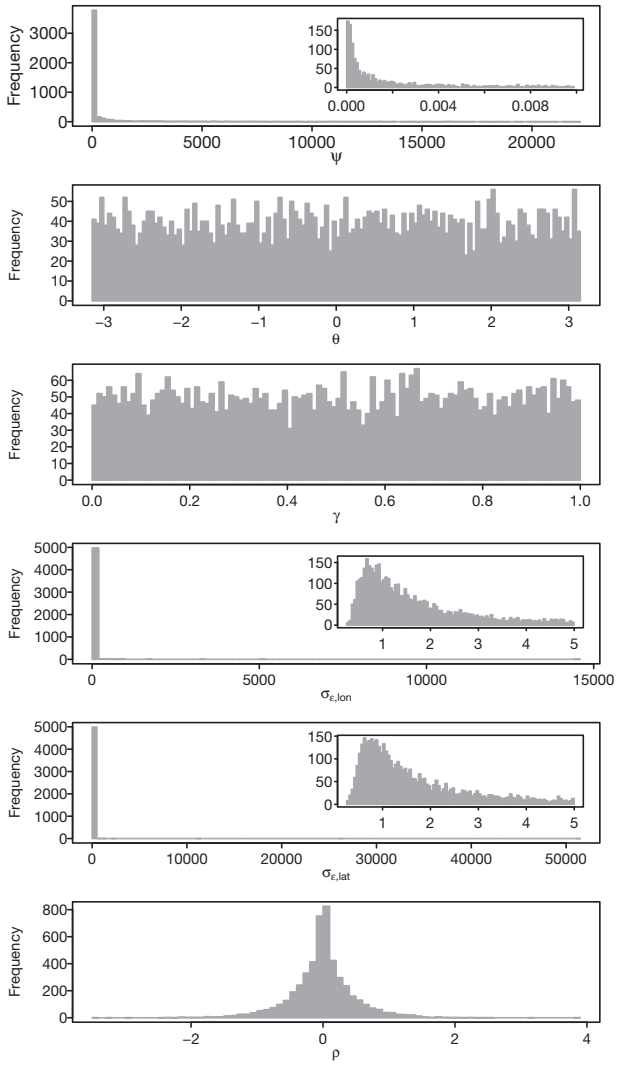

Figure S1.1: Samples from the priors from `bsam`. Histograms of the 5000 values sampled directly from priors. Note that for some parameters a sub-panel displays the lower values.

S1.3 Comparison between the polar bear results from `TMB` and `bsam`

As mentioned in the main text, `TMB` was much faster than `bsam` (Table S1.3). However, `TMB` and `bsam` returned similar parameter estimates and movement tracks. The results from only one movement track is presented in the main text. However, we can see here that the estimated movement tracks with `bsam` and `TMB` were similar and close to the true GPS tracks for the four polar bears (Fig. S1.2). While the state estimates were very close, `TMB` was slightly less accurate than `bsam` (Table S1.3). We can also see that the movement parameters are similar across methods (Table S1.4).

Table S1.3: Comparison of the computational efficiency and accuracy of the two R packages when fitting the DCRW to each polar bear movement path.

ID	Time (min)		RMSE (° lon)		RMSE (° lat)		Distance (km)		Convergence message
	`TMB`	`bsam`	`TMB`	`bsam`	`TMB`	`bsam`	`TMB`	`bsam`	
PB1	0.77	37.69	0.13	0.12	0.03	0.03	3.67	3.60	relative convergence (4)
PB2	2.60	82.70	0.22	0.21	0.05	0.05	5.41	5.22	false convergence (8)
PB3	0.77	38.43	0.32	0.29	0.06	0.06	7.12	6.65	relative convergence (4)
PB4	0.84	36.99	0.25	0.25	0.08	0.08	7.98	7.92	relative convergence (4)

Table S1.4: Comparison of the parameters estimated by `TMB` and `bsam` for each polar bear movement path. The standard errors are presented in brackets.

ID		θ	γ	$\sigma_{\epsilon,lon}$	$\sigma_{\epsilon,lat}$	ρ	ψ
PB1	`TMB`	0.05 (0.03)	0.50 (0.04)	0.34 (0.02)	0.10 (0.01)	-0.05 (0.06)	0.15 (0.01)
PB1	`bsam`	0.05 (0.03)	0.46 (0.04)	0.36 (0.01)	0.12 (<0.01)	-0.03 (0.06)	0.16 (0.01)
PB2	`TMB`	-0.03 (0.02)	0.44 (0.02)	0.39 (<0.01)	0.09 (<0.01)	-0.04 (0.04)	0.14 (<0.01)
PB2	`bsam`	-0.03 (0.03)	0.42 (0.03)	0.39 (0.01)	0.11 (<0.01)	-0.05 (0.05)	0.16 (0.01)
PB3	`TMB`	-0.02 (0.04)	0.42 (0.04)	0.34 (0.02)	0.09 (<0.01)	-0.03 (0.07)	0.11 (0.01)
PB3	`bsam`	-0.01 (0.05)	0.38 (0.04)	0.34 (0.02)	0.11 (0.01)	-0.05 (0.07)	0.12 (0.01)
PB4	`TMB`	-0.07 (0.04)	0.40 (0.03)	0.53 (0.02)	0.15 (<0.01)	-0.18 (0.06)	0.09 (<0.01)
PB4	`bsam`	-0.07 (0.04)	0.38 (0.04)	0.53 (0.02)	0.17 (0.01)	-0.16 (0.06)	0.10 (<0.01)

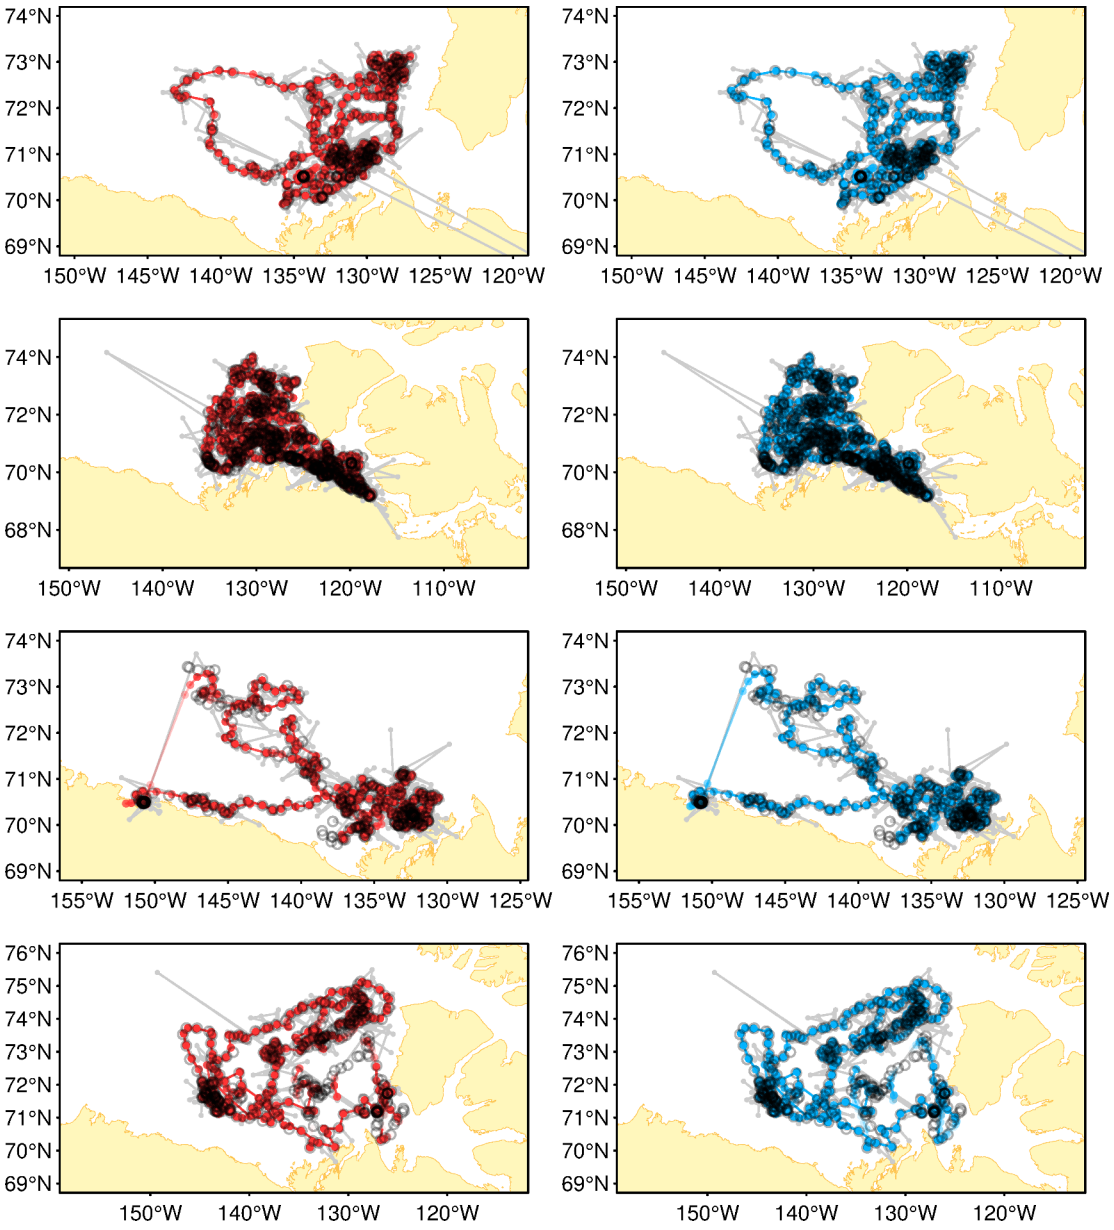

Figure S1.2: Comparison of the estimated paths by `TMB` and `bsam` for four polar bears. The grey points and lines represent the Argos movement path. The colour points and lines represent the estimated movement paths, with red representing the `TMB` estimates and blue the `bsam` estimates. The black circles represent the GPS locations.

S1.4 Likelihood profiles for the DCRW fitted with `TMB`

The likelihood profiles from the parameters from the DCRW indicated that `TMB` had challenges converging and finding the maximum likelihood estimate (MLE). The jagged profiles, as well as the fact that the minimizer could not converge for some fixed parameter values, indicate that the likelihood surfaces were potentially multimodal and difficult to search. Such difficulties could be due to a variety of factors, many of which may be related to discrepancies between the model applied and the data used. See the Discussion section of the main text for further detail.

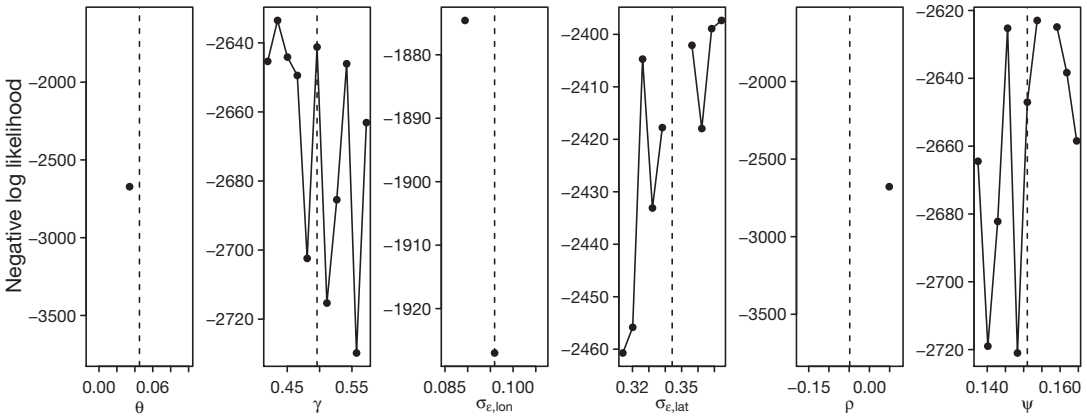

Figure S1.3: Negative log likelihood profile for the parameters estimated with `TMB` for PB1. The points and full lines indicate the estimated negative log likelihood value when the given parameter was fixed but all other parameters were optimized. Discontinuities indicate that the minimizer could not converge at the given parameter value. The dashed line indicates the MLE for the given parameter.

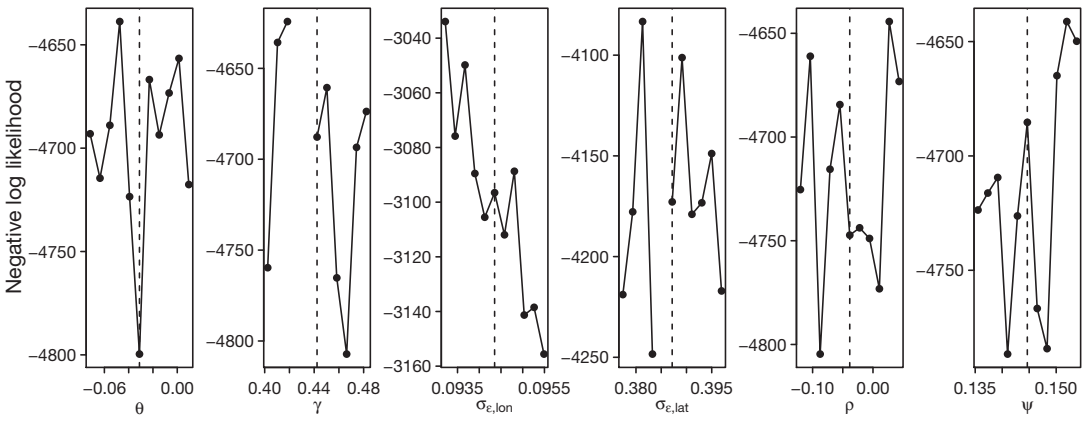

Figure S1.4: Negative log likelihood profile for the parameters estimated with `TMB` for PB2. The points and full lines indicate the estimated negative log likelihood value when the given parameter was fixed but all other parameters were optimized. Discontinuities indicate that the minimizer could not converge at the given parameter value. The dashed line indicates the MLE for the given parameter.

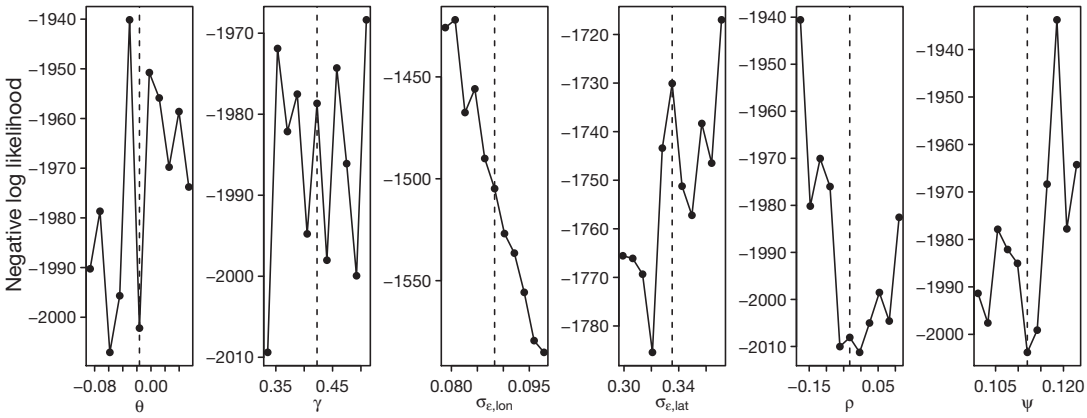

Figure S1.5: Negative log likelihood profile for the parameters estimated with `TMB` for PB3. The points and full lines indicate the estimated negative log likelihood value when the given parameter was fixed but all other parameters were optimized. Discontinuities indicate that the minimizer could not converge at the given parameter value. The dashed line indicates the MLE for the given parameter.

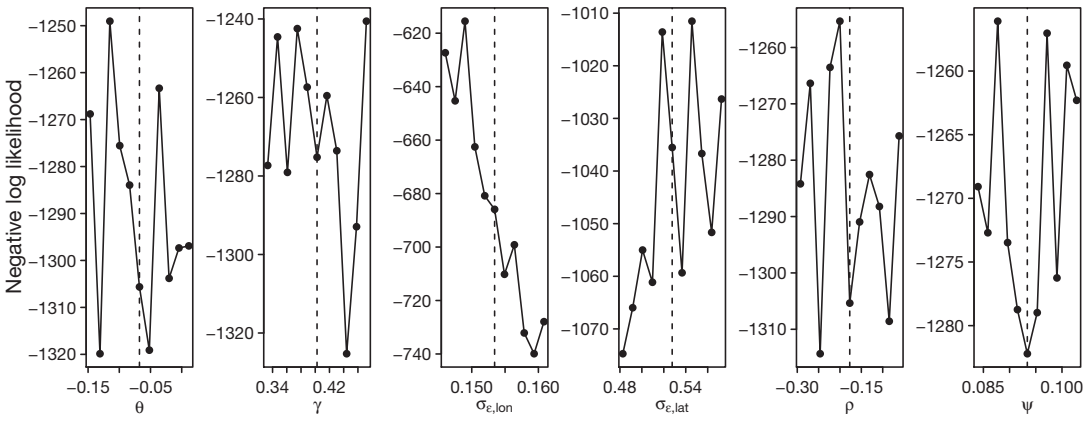

Figure S1.6: Negative log likelihood profile for the parameters estimated with `TMB` for PB4. The points and full lines indicate the estimated negative log likelihood value when the given parameter was fixed but all other parameters were optimized. Discontinuities indicate that the minimizer could not converge at the given parameter value. The dashed line indicates the MLE for the given parameter.

S1.5 Convergence diagnostic for `bsam`

To verify the convergence of the chains when fitting the model through `bsam`, we looked at the diagnostic plots produced by `bsam` and using the potential scale reduction factor. For the latter, values close to one are considered consistent with convergence (Gelman & Rubin 1992). In general, the chains appeared to be well mixed and with low autocorrelation (Figs. S1.7-S1.10). All of the potential scale reduction factors were ≤ 1.02, although a few returned NaN (Table S1.5).

Table S1.5: Potential scale reduction factor for the parameters when fitted with `bsam`.

ID	θ	γ	$\sigma^2_{\epsilon,lon}$	$\rho\sigma_{\epsilon,lon}\sigma_{\epsilon,lat}$	$\sigma^2_{\epsilon,lat}$	ψ
PB1	1.00	1.00	1.00	1.00	1.00	1.01
PB2	1.00	1.01	–	–	–	1.00
PB3	1.02	1.00	–	–	–	1.01
PB4	1.00	1.00	1.01	1.00	1.00	1.00

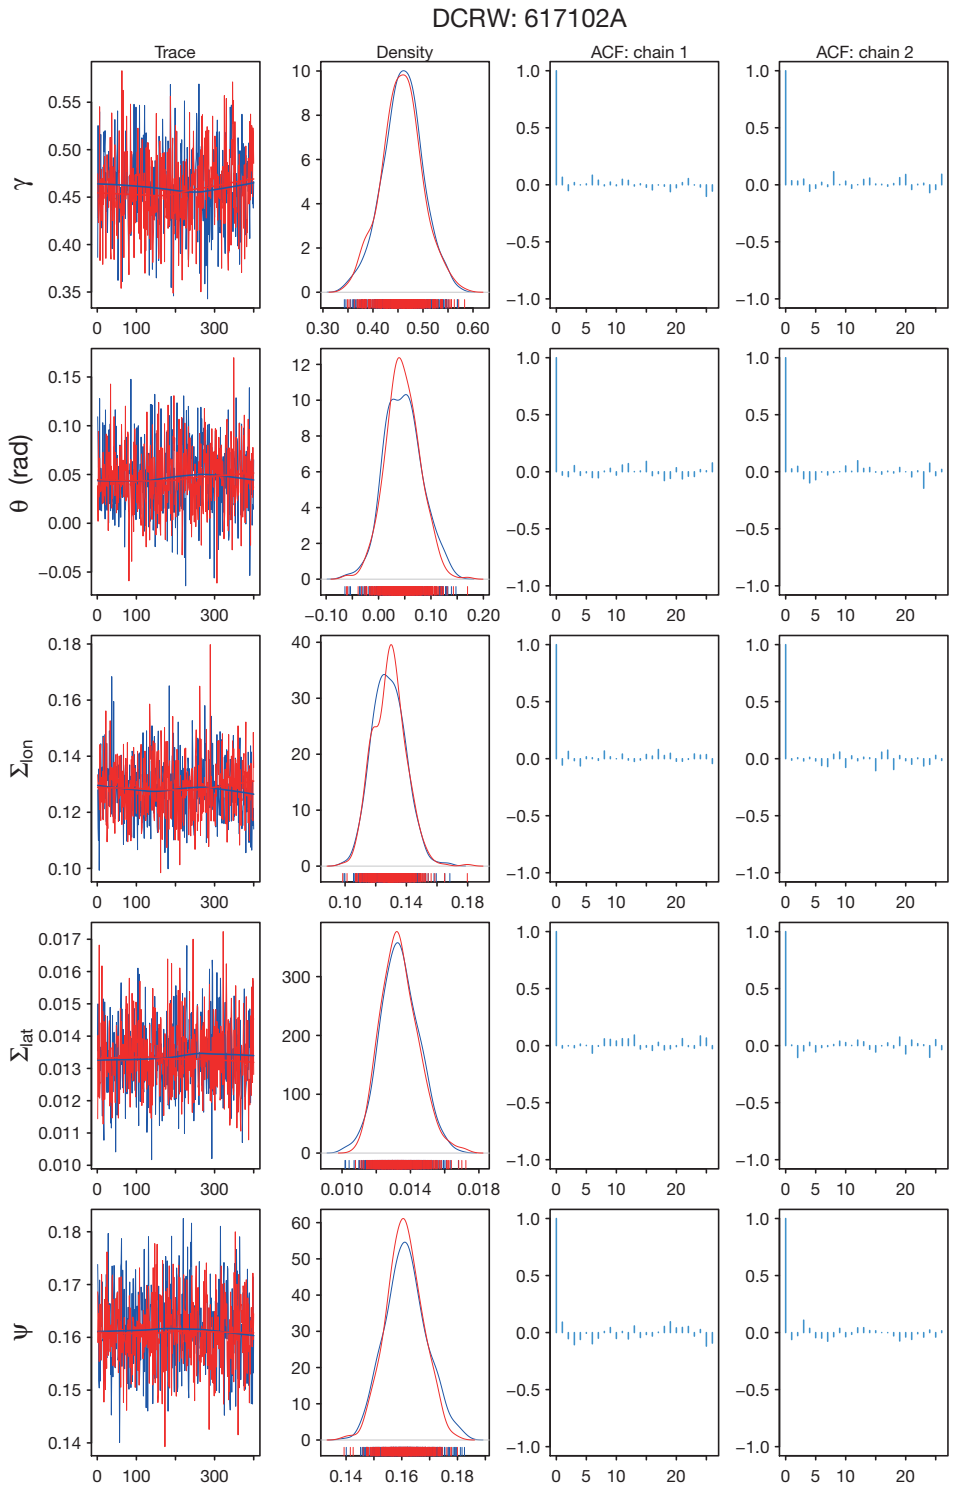

Figure S1.7: Convergence diagnostics from the `bsam` package for PB1.

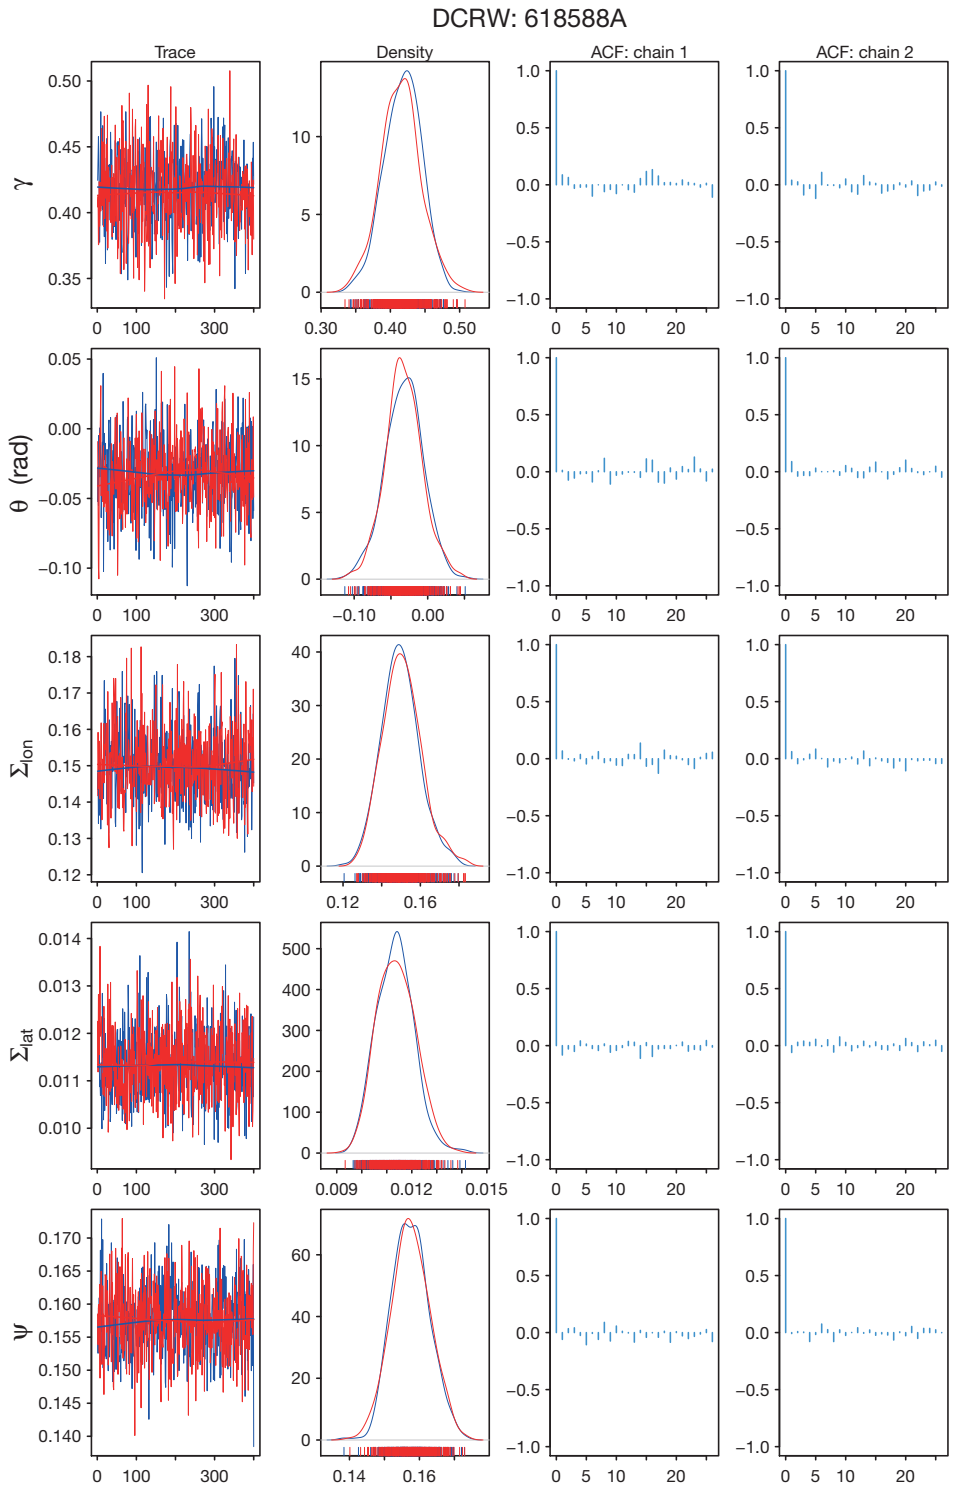

Figure S1.8: Convergence diagnostics from the `bsam` package for PB2.

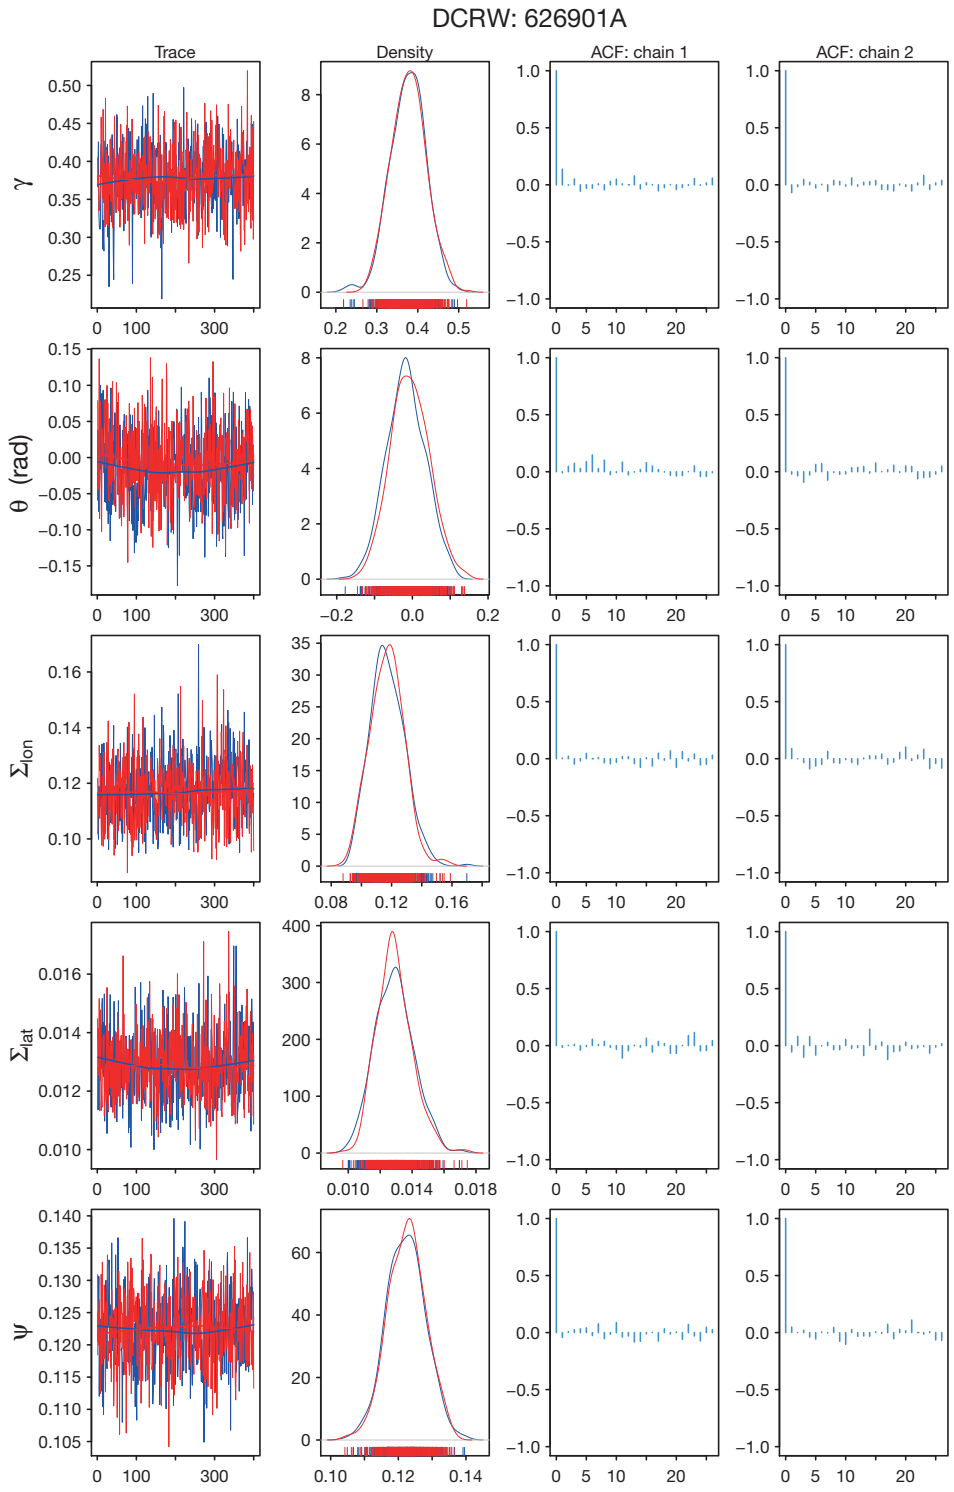

Figure S1.9: Convergence diagnostics from the `bsam` package for PB3.

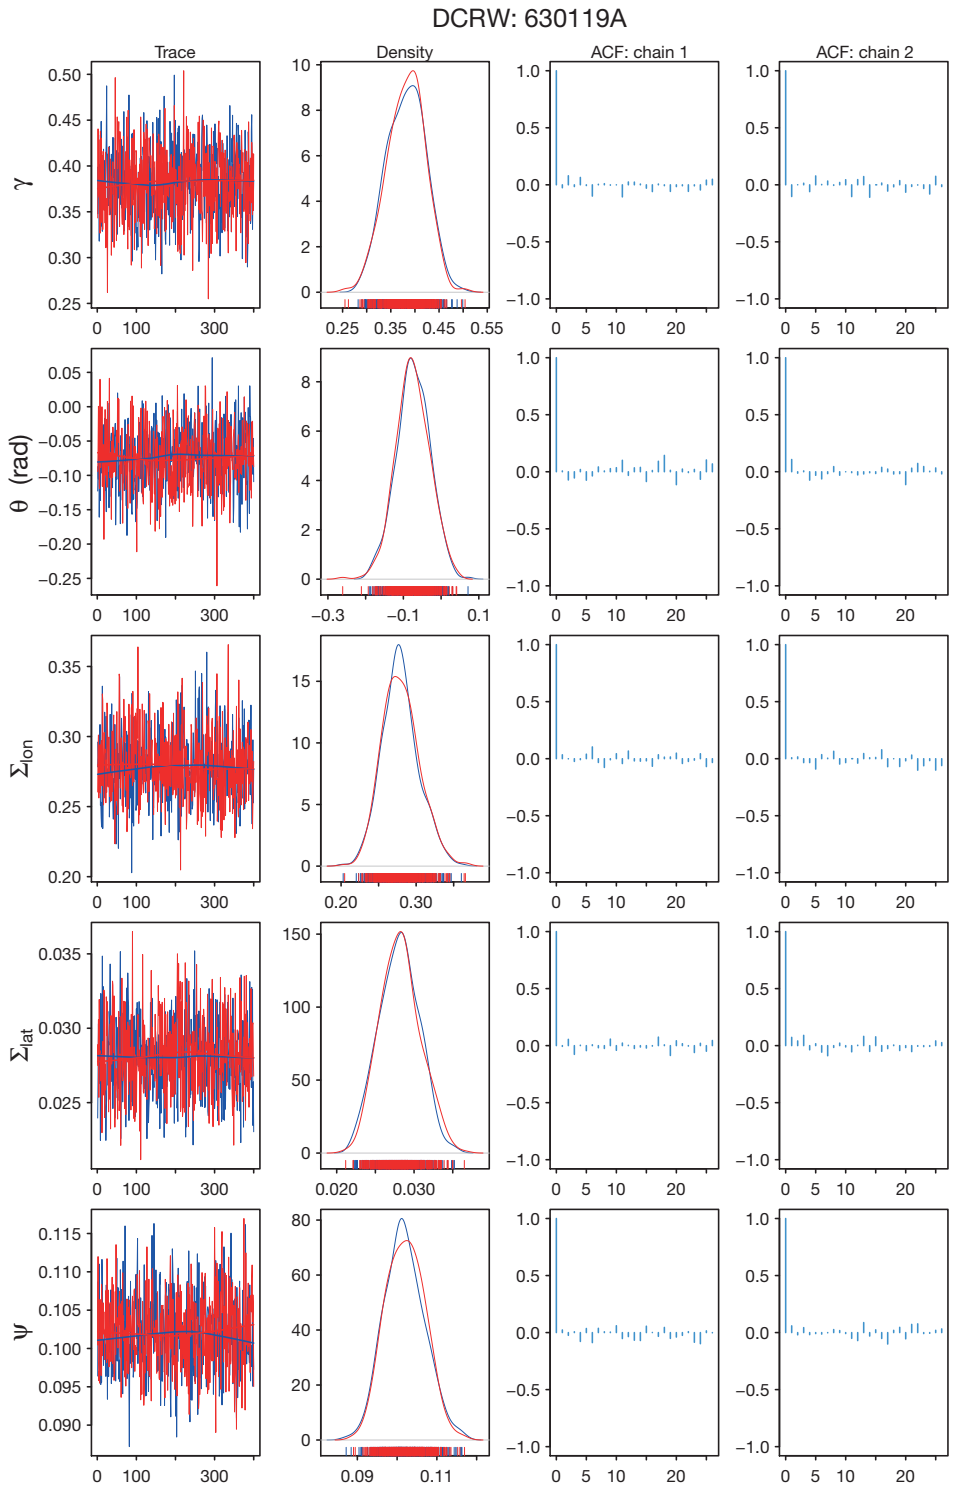

Figure S1.10: Convergence diagnostics from the `bsam` package for PB4.

S1.6 Results from the simulation study

The simulation study showed that `TMB` was much faster than `bsam` (Table S1.6). The accuracy of `TMB` in terms of state estimates was very close to, but on average slightly lower than, that of `bsam` (Table S1.6). The accuracy of `TMB` in terms of estimating the parameters was generally higher than `bsam` for parameters θ, γ, $\sigma_{\epsilon,lon}$, and $\sigma_{\epsilon,lat}$, but lower for ρ and ψ (Table S1.7). For a few simulations, `TMB` could not converge. This is different from false convergence in that `TMB` will not return estimate values. In total, we had $< 4\%$ of the simulations with such problems. We removed these simulations from the analysis. The convergence of the chains when fitting the model through `bsam` was generally satisfactory with a mean potential scale reduction factor of 1.02, 1.01, 1.00 for frequency of 0.5, 1, and 5 locations/day, respectively.

Table S1.6: Comparison of the computational efficiency and state estimate accuracy of the two R packages when fitting the DCRW to simulated movement paths with different data frequency scenarios. For each scenario, we present the mean root mean square error and mean great circle distance across simulations from that scenario. We also present the number of simulations per scenario that returned an error message (i.e. that could not converge at all).

Frequency	Time (min)		RMSE (° lon)		RMSE (° lat)		Distance (km)		N° errors
(loc/day)	`TMB`	`bsam`	`TMB`	`bsam`	`TMB`	`bsam`	`TMB`	`bsam`	
0.5	0.64	15.30	0.57	0.57	0.18	0.18	52.72	52.44	4
1	0.81	17.77	0.37	0.37	0.12	0.12	34.53	34.37	4
5	1.85	35.02	0.04	0.04	0.02	0.02	3.68	3.57	2

Table S1.7: Comparison of the parameter estimate accuracy from the two different R packages when fitting the DCRW to simulated movement paths under different frequency scenarios. For each scenario, we present the root mean square error for the parameters.

Frequency		θ	γ	$\sigma_{\epsilon,lon}$	$\sigma_{\epsilon,lat}$	ρ	ψ
0.5	`TMB`	0.10	0.10	0.06	0.07	0.14	0.11
0.5	`bsam`	0.33	0.33	0.13	0.07	0.09	0.09
1	`TMB`	0.11	0.07	0.03	0.01	0.10	0.04
1	`bsam`	0.15	0.16	0.04	0.04	0.07	0.03
5	`TMB`	0.10	0.04	0.03	0.04	0.06	0.01
5	`bsam`	0.11	0.04	0.01	0.02	0.05	0.01

S2 Additional information for case study 2

We used the second case study to demonstrate how the efficiency of `TMB` and its frequentist framework facilitate model comparison. To do so, we applied two versions of the DCRW to the movement tracks of four rhinoceros auklets. This appendix presents the filter used on the light-based geolocation data and additional analysis results.

S2.1 Effect of filtering on data

As discussed in the main text, we used a filter to remove the most aberrant geolocations. To keep as much information as possible, we filtered each coordinate separately (e.g. we only removed the latitude coordinate when it was clearly aberrant but the longitude estimate appeared adequate). To do so, we applied a two-step filter to each coordinate time-series. First, we removed geolocations outside the possible range of these individuals (latitude $\leq 10°$N, longitude $\leq 110°$W). Second, we removed any extreme values that disrupted the continuity of the path (i.e. estimates > 1400 km from the average coordinate value within a nine-day window). The distance was calculated using the great-circle distance using `distVincentyEllipsoid` function from the R package `geosphere`. Given that rhinoceros auklets are estimated to move on average 240 km/day during the non-breeding period (based on the 10 km/hr estimate of Takahashi et al, 2015), a threshold of 1400 km would keep many potentially unrealistic locations. Around the equinoxes (September 23, March 20), we used a lower threshold of 1000 km. Similar to Ballard et al (2010), we accounted for the differential error levels before and after the two equinoxes (equinox periods used: September 13 to October 14 and February 27 to March 30). Because we used the error estimates provided by the tag, we additionally removed geolocations for which the tag error value was missing. Fig S2.1 demonstrates the data points that were removed.

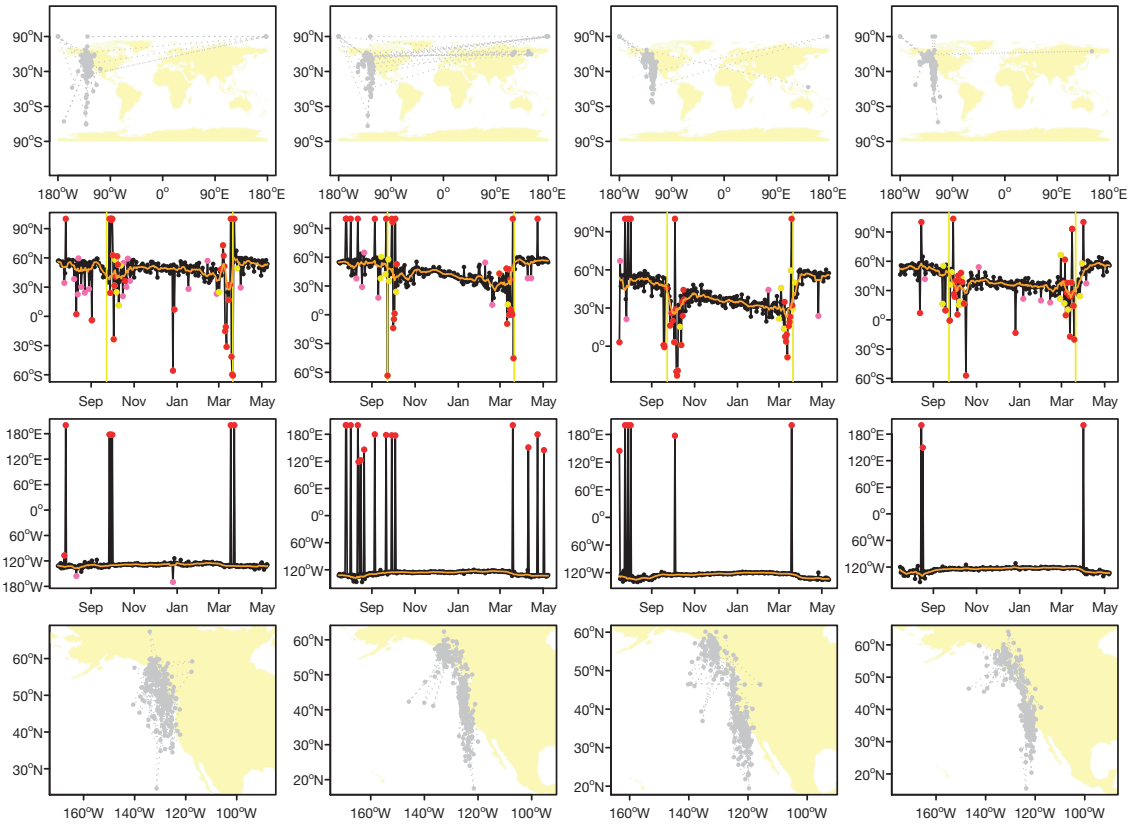

Figure S2.1: The effect of the filtering method on the data. The first row represents the raw data, when the coordinates outside the coordinate system are set to either boundary. The second and third rows show the raw data with the removed values in colour. Red represents the coordinate values that were either outside the coordinate system, outside the possible range for these individuals, or for which the tags did not provide an error estimate. Pink represents the values that are 1400 km from the mean (shown in orange), and yellow represents the values that are 1000 km from the mean while during the equinox period. The yellow lines represent the equinoxes. The fourth row represents the data when the aberrant geolocations are removed. Note that only the locations for which we have both coordinates are displayed.

S2.2 Parameter and state estimates for the rhinoceros auklet movement tracks

Table S2.1 presents the parameter estimates for the best version (M_T) when applied to each of the four rhinoceros auklet movement tracks. In addition, Fig. S2.2 presents the movement tracks returned by the two versions of the model for each of the four rhinoceros auklets.

Table S2.1: Parameter estimates for the four rhinoceros auklet movement paths for the best version (M_T, the version with t-distributed measurement errors). The standard error of the parameters is displayed in brackets.

Parameters (units)	Auklet 1	Auklet 2	Auklet 3	Auklet 4
γ	0.29 (0.07)	0.81 (0.07)	0.61 (0.17)	< 0.01 (< 0.01)
$\sigma_{\epsilon,lon}$ (°)	0.37 (0.06)	0.12 (0.03)	0.21 (0.08)	0.47 (0.06)
$\sigma_{\epsilon,lat}$ (°)	0.83 (0.06)	0.31 (0.09)	0.62 (0.21)	1.31 (0.20)
α_{lon}	0.28 (0.02)	0.33 (0.02)	0.30 (0.02)	0.36 (0.03)
α_{lat}	0.27 (< 0.01)	0.33 (0.03)	0.37 (0.04)	0.33 (0.03)
df_{lon}	3.34 (0.28)	3.00 (< 0.01)	3.00 (< 0.01)	3.00 (< 0.01)
df_{lat}	3.62 (0.05)	3.61 (0.96)	5.90 (2.61)	3.00 (< 0.01)

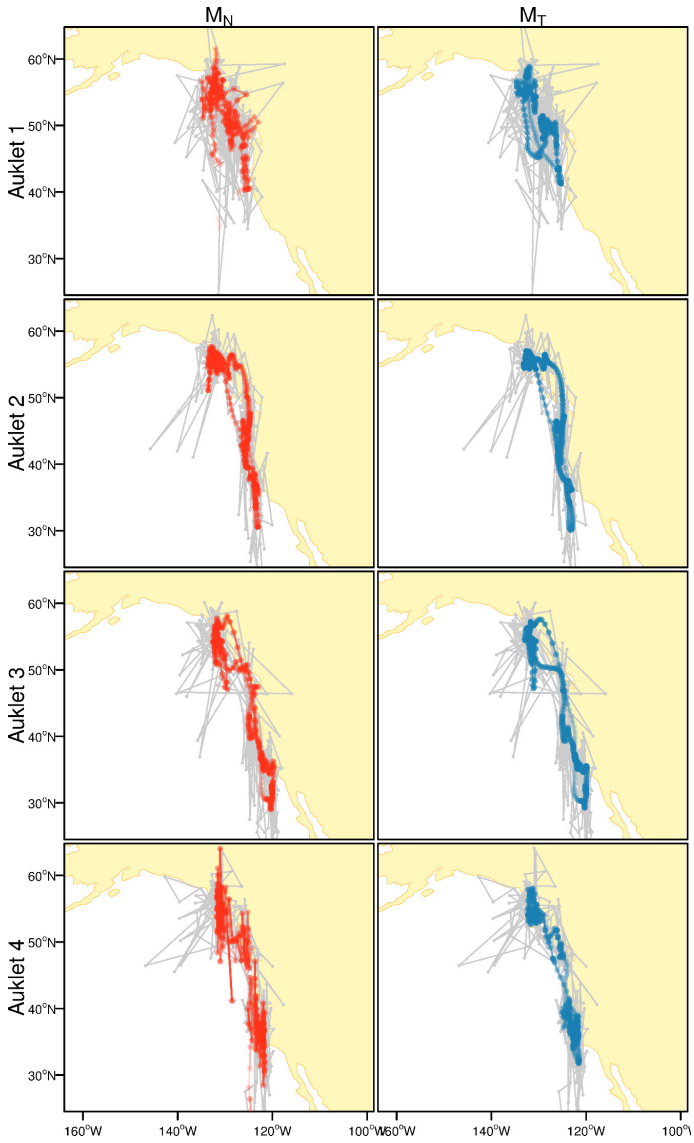

Figure S2.2: Comparison of the fit to each individual of the two model versions: M_N, the version with normally distributed measurement errors; and M_T, the version with t-distributed measurement errors. The grey points and lines represent the observed data and the colour points and lines the estimated movement path. The intensity of the colour represents the confidence in the location estimate and any location with $(S.E.(x_{lat,t}) + S.E.(x_{lon,t})) \geq 5°$ is not displayed.

S2.3 Simulation results from case study 2

Table S2.2: Number of model versions, for each simulation scenario, selected as best according to AIC. For each version simulated (rows), the number of simulations for which the given fitted version (columns) had the lowest AIC is shown.

	Fitted version	
Simulated version	M_N	M_T
M_N	97	3
M_T	0	100

S3 Additional information for case study 3

This appendix presents additional information on the Fastloc-GPS data of the grey seals. In particular, Table S3.3 presents the length of each time series and additional information on the time difference between locations.

Table S3.3: Extra information on the four movement tracks of the grey seals. The total % of locations removed by the quality filter (any location with < 5 satellites or tag residuals > 35) are presented. In addition, the % of locations missing based on the number of locations expected by 15 min scheduled are presented. The maximum and median time intervals between locations are presented.

ID	n	% missing		Δt_i (min)		Start date	End date
		filter	total	max	median		
Seal 1	10,332	6	46	54h 20	18	2013-06-23	2014-01-09
Seal 2	8,085	22	53	55h 48	18	2013-06-25	2013-12-22
Seal 3	7,395	23	57	63h 42	20	2013-06-26	2013-12-22
Seal 4	11,829	8	36	40h 43	18	2013-06-27	2014-01-06

Bibliography

Ballard G, Toniolo V, Ainley DG, Parkinson CL, Arrigo KR, Trathan PN (2010) Responding to climate change: Adélie Penguins confront astronomical and ocean boundaries. Ecology 91:2056–2069

Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences. Stat Sci 7:457–511.

Takahashi A, Ito M, Suzuki Y, Watanuki Y, Thiebot J, Yamamoto T, Iida T, Trathan P, Niizuma Y, Kuwae T (2015) Migratory movements of rhinoceros auklets in the northwestern Pacific: connecting seasonal productivities. Mar Ecol Prog Ser 525:229–243