Global Biogeochemical Cycles

Supporting Information for

Impact of Lagrangian Sea Surface Temperature Variability on Southern Ocean Phytoplankton Community Growth Rates

Jessica Zaiss[1], Philip W. Boyd[2], Scott C. Doney[3], Jon N. Havenhand[4], Naomi M. Levine[5]

[1]Department of Earth Science, University of Southern California

[2]Institute for Marine and Antarctic Studies, University of Tasmania

[3]Department of Environmental Sciences, University of Virginia

[4]Department of Marine Science, University of Gothenburg

[5]Department of Marine and Environmental Biology, University of Southern California

Contents of this file

Introduction

This supplemental material contains the results of the sensitivity analyses we performed on thermal niche width, final SST in the idealized simulations, and the magnitude of the imposed minimum biomass, as well as the statistical analysis to determine the significance, or lack of, the differences between the sensitivity tests and the results in the main text. Also included are figures to supplement the findings in the main text such as results for the broad shaped reaction norms and the decrease ΔSST results for the skewed reaction norms.

S.1 Impact of final temperature in the idealized simulations

To assess the sensitivity of our choice of final SST of 15^{o}C for the idealized simulations, we performed 100 idealized simulations with final SSTs of 10^{o}C and 20^{o}C with the same rates and magnitudes of temperature change as presented in the main text. Specifically, we compared the percent difference between the individual-based model and the Q_{10} parameterization, relative to Q_{10}, at the timestep when SST stabilizes as well as the length of time it took for the community growth rates to equilibrate to steady conditions, referred to as the memory length in the main text.

The final SST of the idealized profiles did not impact the results of our study. The offset between the Q_{10} parameterization and the individual-based model remained statistically similar (95% CI, see section S.7 for description of statistical analyses) for all the simulations (Figure S1). Additionally, because the results are presented in terms of generations rather than absolute time, the length of the memory effect is also statistically the same (95% CI) for all simulations (Figure S2).

S.2 Impact of thermal niche width in idealized simulations

The width of the thermal niche had a varying, but predictable, impact on the percent difference from Q_{10}. To test the impact of the thermal niche width, we ran 100 simulations of the same idealized SST profiles as in the main text with both narrower and wider thermal niches for the individuals. Making the thermal niche more narrow (10.5^{o}C) relative to the simulations in the main text (thermal niche = 14^{o}C) did not have a significant (95% CI) impact on the percent difference between the individual-based model and the Q_{10} growth parameterization, for either shape of reaction norm (Figure S3, left column). Increasing the thermal niche from 14^{o}C to 20.5^{o}C (Figure S3, right column) also did not have a significant impact on the offset from the Q_{10} parameterization for skewed reaction norms (95% CI) but did for the broad reaction norms (95% CI). A wider thermal niche for the broad reaction norms decreased the percent difference from Q_{10} by an average of 10.1% with simulations in which ΔSST changed over 7 days experiencing the largest decreases of up to 29.4%.

The width of the thermal niche in conjunction with the magnitude of SST change impacted the memory length. For small (2-3^{o}C) and large (8-9^{o}C) SST changes, wider thermal niches produced shorter memory effects (95% CI) by an average of 2.1 generations ± 3.4 generations (1σ) for the broad reaction norms and 0.6 generations ± 1.6 generations (1σ) for the skewed reaction norms. Conversely, wider thermal niches that experienced moderate SST changes (4-7^{o}C) had longer memory effects than the default thermal niches by 1.8 generations ± 4.0 generations (1σ) for the broad and 0.5 generations ± 1.0 generations (1σ) for the skewed reaction norms. This was seen across both sets of simulations for the broad reaction norms (Figure S4, bottom row). For skewed reaction norms, decreasing the thermal niche width relative to the default width did not have a significant impact on the memory length (95% CI) (Figure S4, top row). Regardless of thermal niche width, the overall relationship between memory length and the rate and magnitude of SST change was the same as the simulations in the main text and did not change our results or conclusions.

For broad shaped reaction norms, wider thermal niches mean that individuals were able to continue to grow over a larger span of SSTs on either side of their optimum growth temperature (T_{opt}) compared to individuals with narrower thermal niches. When temperature changes were small (2-3^{o}C), the biomass weighted community growth rate was able to better track small changes in SST because the SST did not go outside of the thermal niche. Similarly, for large SST changes (8-9^{o}C), the community was able to respond to the SST changes more quickly than a community with a narrower thermal niche because more individuals were able to grow at the final SST. When SST changes were more moderate (4-6^{o}C), the individuals in the original environment could continue to grow over a larger range of temperatures past their T_{opt} which meant that those best suited for the new environment had more biomass to overcome before making a significant contribution to the biomass-weighted community growth rates compared to individuals in a community with narrow thermal niches.

The memory length for the skewed reaction norms was less affected by the width of the thermal niche due to the asymmetry of the reaction norm shape. By increasing the width the reaction norm, but keeping the maximum growth rate and T_{opt} the same, the part of the reaction norm that was extended corresponded to relatively low growth rates. So even though individuals could grow at a larger temperature range, that growth did not have a large impact on the memory length.

The impact of thermal niche width on the difference from Q_{10} and the length of the memory effect did not change any of the conclusions of the manuscript. Across all the simulations, larger and faster SST changes resulted in the largest offsets from Q_{10} and moderate SST changes induced the longest memory effects.

S.3 Model sensitivity to minimum biomass parameter

In the main text, we imposed a minimum biomass of 0.001 mmol C m^{-3} such that no individual was allowed to go extinct, akin to the "everything is everywhere" principle (Hutchinson, 1961). To test the sensitivity of our results to this parameter, we ran 100 simulations with the same idealized SST profiles with a minimum biomass of 0.0001 mmol C m^{-3}, an order of magnitude smaller. For both the skewed shaped and broad shaped reaction norms, lower minimum biomass generally increased both the offset from Q_{10} simulated growth rates (95% CI, Figure S5) and the memory length (95% CI, Figure S6). The difference from Q_{10} increased by an average of 1.5% ±8.6% (1σ) for the broad and 2.6% ± 8.9% (1σ) for skewed shaped reaction norms, but ranged as high as 31.7% (broad) and 27.3% (skewed). For small ΔSSTs (2-3^{o}C), lower minimum biomass slightly decreased the difference between Q_{10} simulated growth but as ΔSSTs increased, so did the offset. Memory lengths increased by an average of 4.0 ± 4.1 (1σ) generations for the broad reaction norms and 3.0 ± 3 (1σ) generations for the skewed reaction norms, but ranged as high as 12.6 generations (broad) and 10.6 generations (skewed) longer for the smaller minimum biomass. A lower minimum biomass meant that individuals with the minimum biomass contributed less to the overall biomass-weighted community growth rate, resulting in lower growth rates and larger departures from Q_{10}. This also meant that those individuals best suited for the new environment started growing with lower biomass and thus took longer to overcome the previously accumulated biomass from the initial conditions which resulted in longer memory lengths. As such, the results presented in the main text are a conservative estimate of the difference from Q_{10} and memory length.

The overall patterns remained the same between both minimum biomass simulations. The direction of the ΔSST change did not impact the memory length for the broad reaction norms whereas decreasing ΔSSTs yielded longer memory lengths for the skewed reaction norms for both sets of simulations. In both sets of simulations, the moderate ΔSSTs resulted in the longest memory lengths.

S.4 Comparison of Ecosystem Model Choice

We compared the community growth rates from several different models to ensure that the results we found were not the result of our choice of model. We found that all models showed similar responses in community growth rate. Below is a description of each of the models used in this comparison.

The biomass of each individual (P_i, in mmol C m^{-3}) was calculated as

$$\frac{dP_i}{dt} = \mu_i * P_i - loss \qquad \text{Eq. S1}$$

where $\mu_{i,t}$ (day $^{-1}$) is the individual growth rate at time *t*. Here we investigated different formulations for the loss term.

Linear Mortality

We started with simple linear mortality, where loss scales linearly with biomass, similar to Moisan et al. (2002).

$$\frac{dP_i}{dt} = \mu_i * P_i - m * P_i \qquad \text{Eq. S2}$$

We found that, the mortality had to be set to unrealistic values (approx. equal to Q_{10} values) in order to keep biomass from exponentially increasing. However, this model does still show a dip in community growth rates with changes in SST that is described in the main text.

Quadratic Mortality (used in the main text)

A more common approach is to represent loss as a quadratic mortality:

$$\frac{dP_i}{dt} = \mu_i * P_i - m * {P_i}^2 \qquad \text{Eq. S3}$$

Simulating phytoplankton loss as quadratic mortality showed the same dip in community growth rates as SSTs begin to as described in the main text. The overall magnitude of the loss term is consistent with the other models also.

Simple Ecosystem

We also tested a more complex ecosystem model with linear mortality and loss due to grazing.

$$\frac{dP_i}{dt} = (\mu_i - m) * P_i - g * \frac{P_i}{P} * Z * P_i \qquad \text{Eq. S4}$$

where g is the temperature dependent grazing (m^3 mmol C^{-1} day $^{-1}$) and Z is the total zooplankton biomass (mmol C m^{-3}). To keep our phytoplankton and zooplankton growth internally consistent, we simultaneously solve for the change in total phytoplankton biomass (P) and zooplankton biomass (Z) over time (where $P = \sum P_i$ for individuals whose biomass is greater than the minimum) using the following equations:

$$\frac{dP}{dt} = \lambda * P - g * Z * P \qquad \text{Eq. S5}$$

$$\frac{dZ}{dt} = 0.3 * g * Z * P - m_z * Z \qquad \text{Eq. S6}$$

where λ (day^{-1}) is the biomass weighted community growth rate from all P_i > minimum biomass, 0.3 is the zooplankton efficiency, and m_z is the zooplankton mortality rate (day^{-1}). Re-solving for total P instead of using the sum of the individual biomasses allowed us to avoid issues with resetting low biomass individuals to the minimum biomass which constantly adds biomass to the system. This resulted in predator-prey oscillations (Figure S7) but also showed the dip in community growth rates as SSTs began to change.

Constant grazing

This model followed the same equations outlined for the Simple Ecosystem model above, but instead of solving for how zooplankton biomass changes over time, we calculate Z for each timestep as:

$$Z == -0.0187 * \lambda + 5 * \frac{\lambda}{a} \qquad \text{Eq. S7}$$

where λ (day $^{-1}$) is the community growth rate defined in the main text (Equation 5) and *a* (day $^{-1}$) is the growth rate from the Q_{10} parameter (Equation 2). This formulation provided a relatively constant grazing pressure which prevented predator-prey oscillations. As seen with the other formulations, this resulted in a decrease in community growth rates as SSTs change.

S.5 Statistics Calculations for Sensitivity Tests

To calculate the potential significance of results from the sensitivity tests, we performed Type II linear regression and tested the significance of the slope against a value of 1. The regression was performed using the *lsqfitma* function in Matlab made available from the Monterey Bay Aquarium Research Institute (https://www.mbari.org/index-of-downloadable-files/). This provided a slope and the uncertainty on that slope. Using these data, we then calculated the Z test statistic as:

$$Z = \frac{x - \mu}{\sigma\sqrt{\frac{1}{N}}} \qquad \text{Eq. S8}$$

where x is the slope to test against, here set to one, μ is the slope from the Type II regression, σ is the standard deviation on the slope, and N is the number of independent tests to find μ, which is one for the *lsafitma* regression. Once Z is calculated, we compare this to the standard score based on a 95% confidence interval which corresponds to a standard score of ±1.96. If Z is outside of this range, we reject the null hypothesis that the slope, μ is equal to one. Otherwise, we fail to reject the null hypothesis.

S.7 Acclimation Rates

To test the impact of different acclimation timescales we performed sensitivity tests in which we imposed systematic acclimation rates for all phenotypes in the model ranging from 0.2 ^{o}C day^{-1} to 0.6 ^{o}C day^{-1} in increments of 0.1 ^{o}C day^{-1}. These rates are consistent with acclimation rates determined for the Southern Ocean diatom *F. cylindrus* (Robert Strzepek, personal communication). We then ran the model with the idealized SST profiles for a ΔSST = 2^{o}C in 7 days (0.29 ^{o}C day^{-1}), 3^{o}C in 7 days (0.43 ^{o}C day^{-1}), 4^{o}C in 7 days (0.57 ^{o}C day^{-1}) and 5^{o}C in 21 days (0.24 ^{o}C day^{-1}). These intervals corresponded to the magnitudes and rates of change most commonly experienced by the drifter trajectories (see Section 3.1) for which the rate of change was greater than 0.2 ^{o}C day^{-1}.

Acclimation in the model was represented as a linear rate of change with the growth rate following the thermal reaction curve. Specifically, if SST rapidly changed from 15^{o}C to 16^{o}C in one day, a phenotype with an acclimation timescale of 0.2^{o}C day^{-1} would move from the growth rate at 15^{o}C to the growth rate at 15.2^{o}C. If the SST then held constant at 16^{o}C, the phenotype would acclimate by the end of the fifth day.

Table S1. Table S1. Results of SST_{max} variability analysis.

	7 days	21 days	45 days	90 days
# of data points	729,262	465,785	273,997	1593
Mean ΔSST_{max}, ^{o}C	0.9	1.7	2.7	4.2
Standard Deviation, ^{o}C	0.7	1.0	1.5	2.0
Median ΔSST_{max}, ^{o}C	0.7	1.5	2.4	3.9
Mode ΔSST_{max}, ^{o}C	0.4	1.1	2.0	2.0
Skewness ΔSST_{max}	2.5	1.6	0.9	0.9

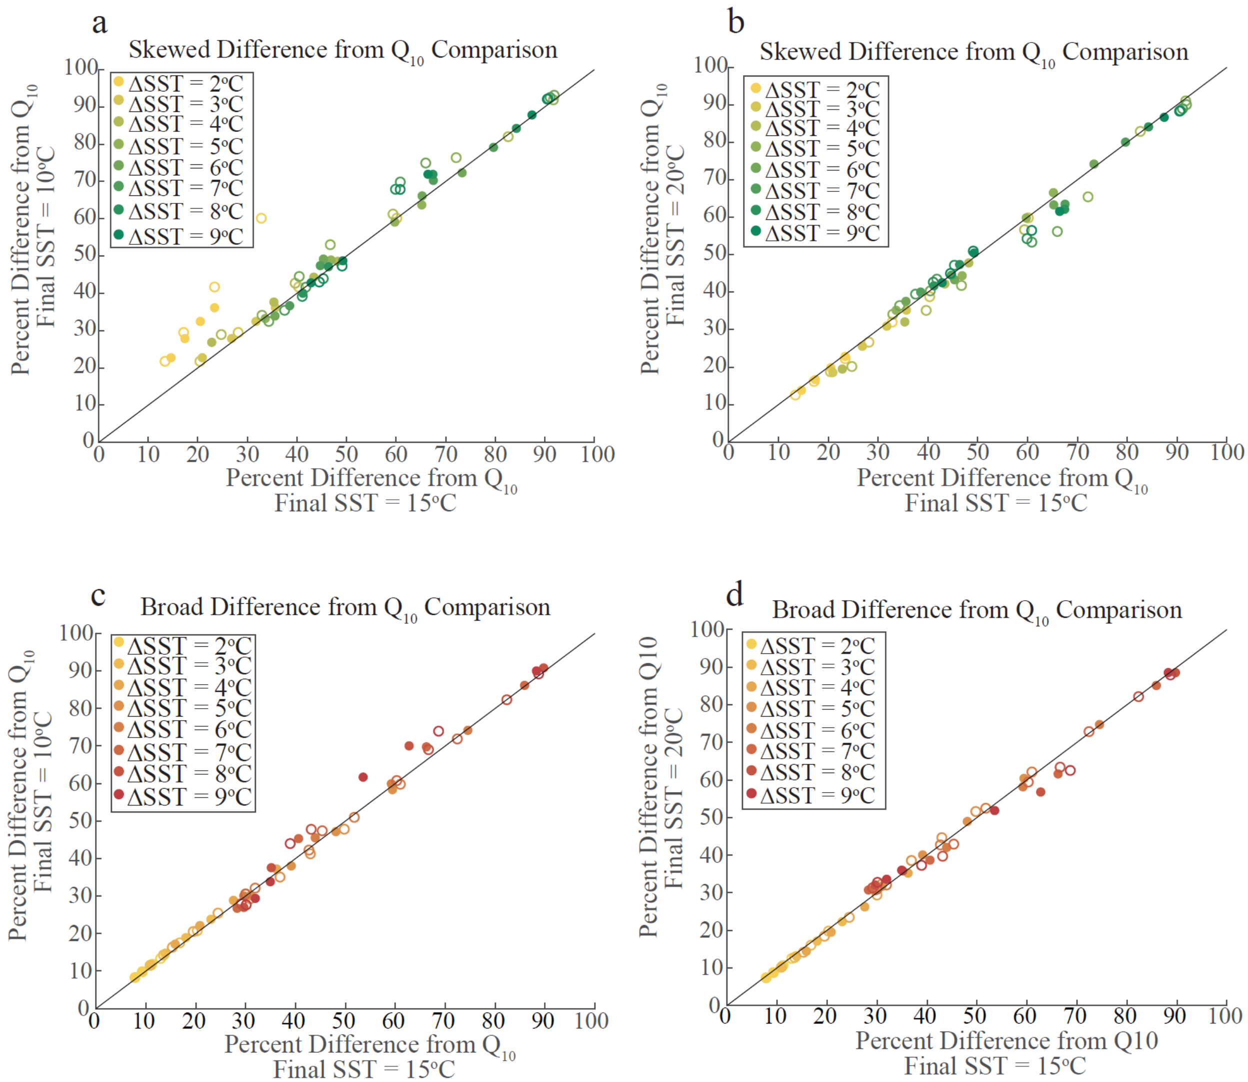

Figure S1. The impact of final SST on percent difference between the individual based model and the Q_{10} parameterization, relative to Q_{10}. The results from the simulations in the main text are compared to simulations with final SSTs of 10ºC (a,c) and 20ºC (b,d) for both the skewed (top row) and broad (bottom row) shaped reaction norms. Open data points represent decreasing ΔSSTs and filled in data are increasing ΔSSTs. The black line indicates the 1-1 line. There is no statistical difference between simulations with differing final SSTs (95% CI).

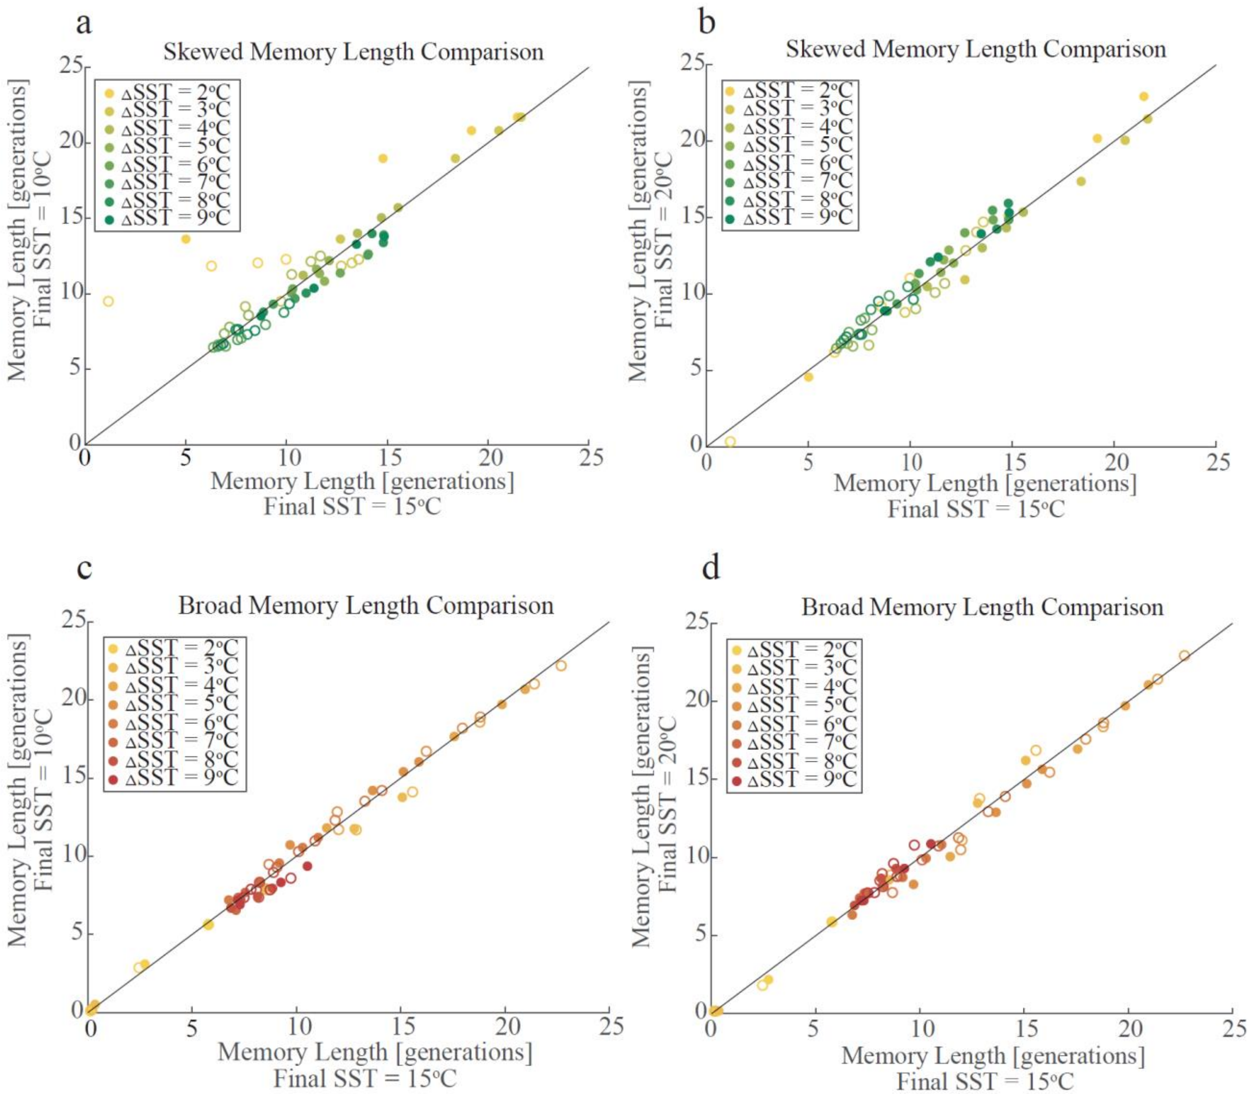

Figure S2. The impact of final SST on memory length. The results from the simulations in the main text are compared to simulations with final SSTs of 10ºC (a,c) and 20ºC (b,d) for both the skewed (top row) and broad (bottom row) shaped reaction norms. Open data points represent decreasing ΔSSTs and filled in data are increasing ΔSSTs. The black line is the 1-1 line. There is no statistical difference between simulations with differing final SSTs (95% CI).

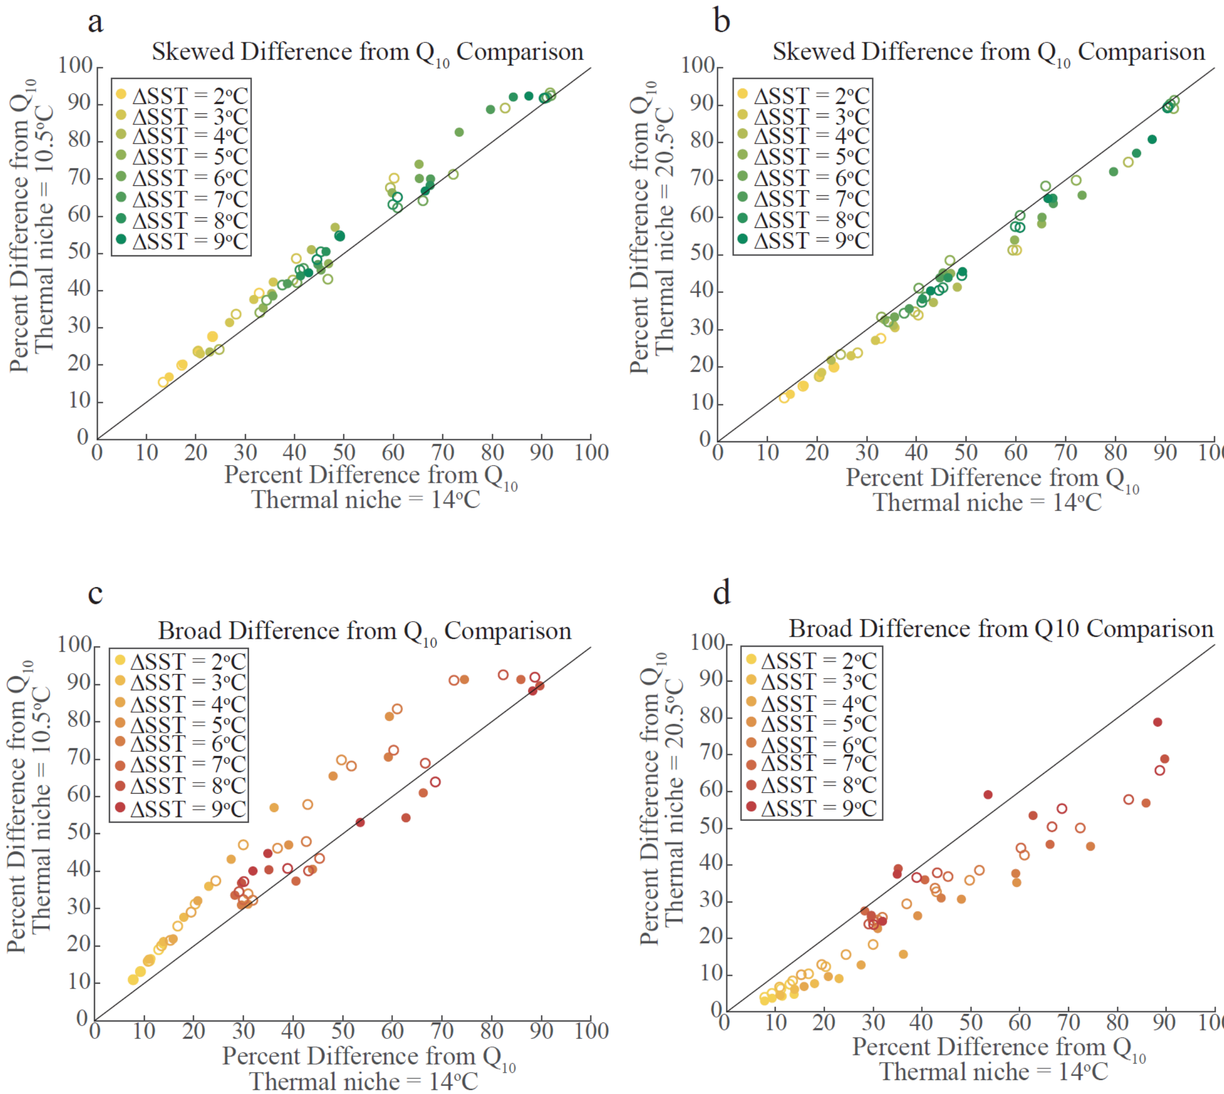

Figure S3. The impact of thermal niche width on percent difference from Q_{10}. The results from the simulations in the main text are compared to simulations with narrower (a,c) and wider (b,d) for both the skewed (top row) and broad (bottom row) shaped reaction norms. The black line is the 1-1 line. Closed data points represent increasing ΔSSTs and open circles represent decreasing ΔSSTs. For broad reaction norms, increasing the thermal niche increases the difference from the Q_{10} parameterized growth rates. There was no significant difference between the simulations for skewed reaction norms.

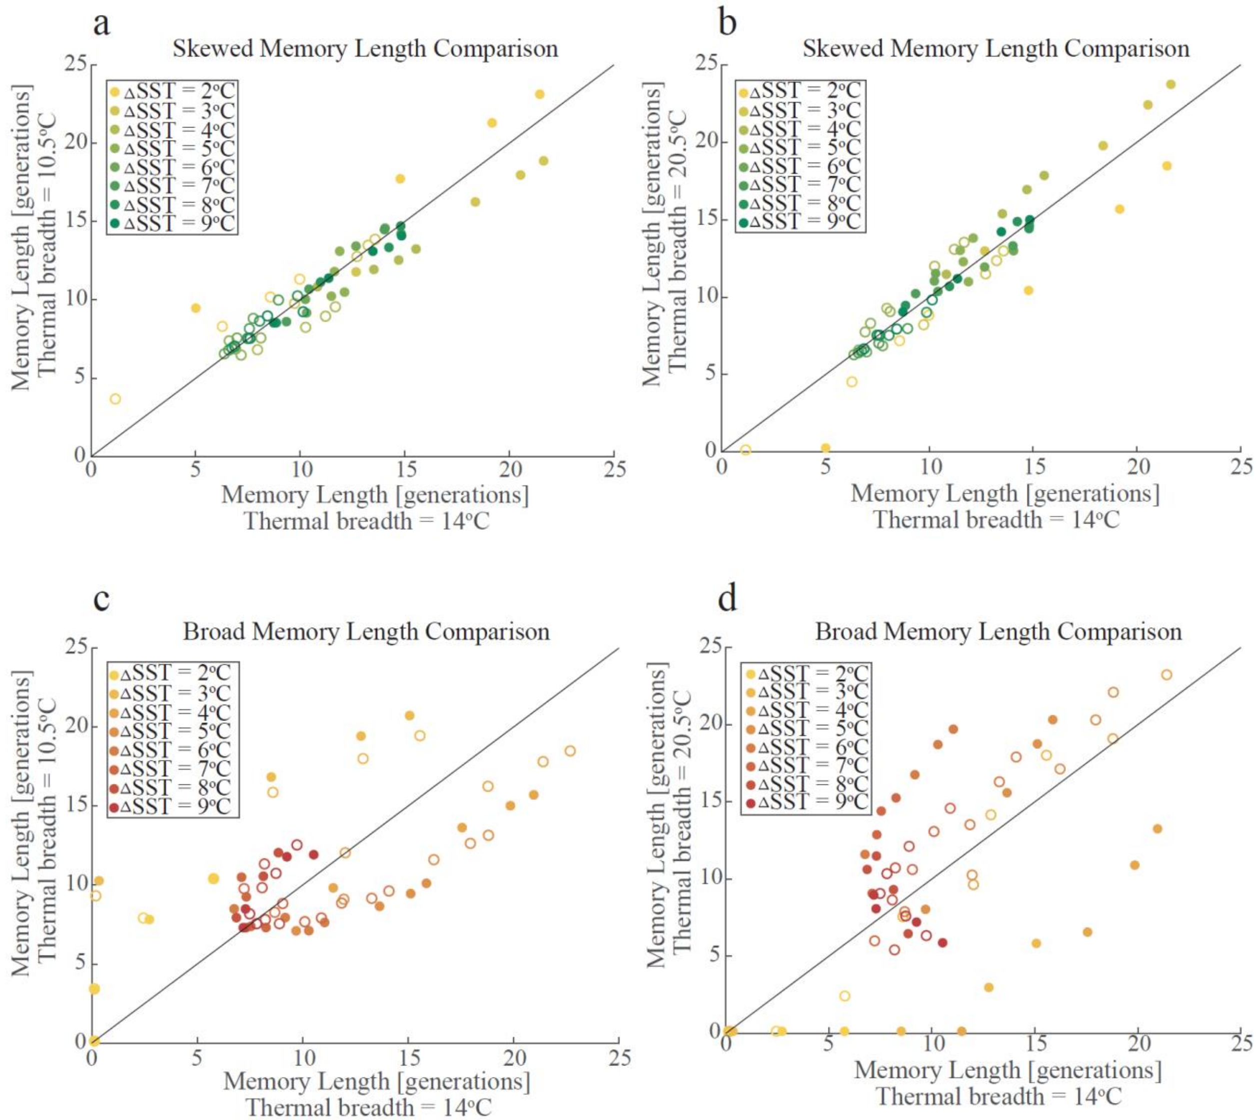

Figure S4. The impact of thermal niche width on memory length. The results from the simulations in the main text are compared to simulations with narrower (left column) and wider (right column) for both the skewed (top row) and broad (bottom row) shaped reaction norms. The black line is the 1-1 line. Closed data points represent increasing ΔSSTs and open circles represent decreasing ΔSSTs. Broad reaction

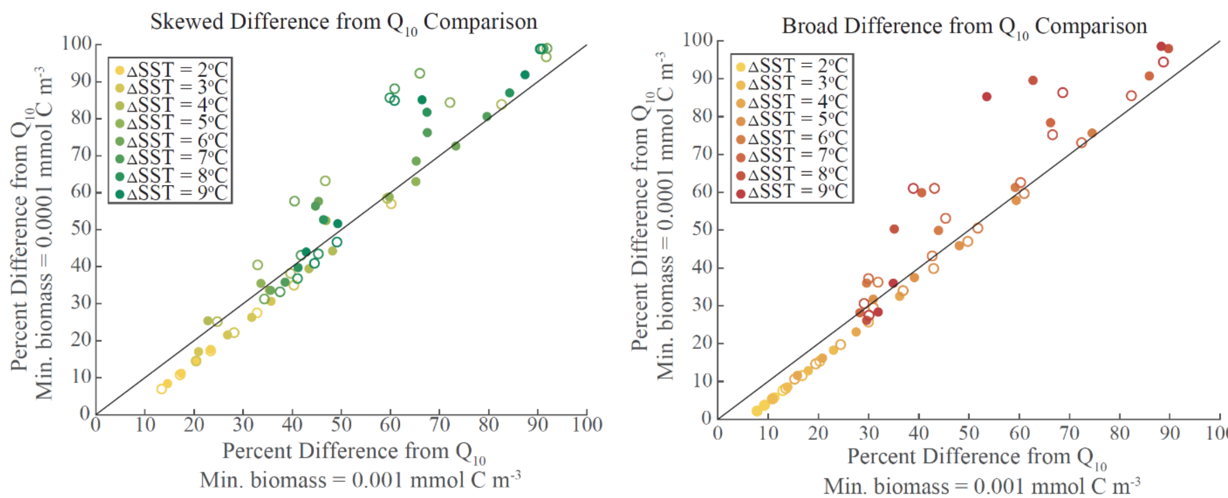

Figure S5. The impact of minimum biomass on deviation from Q_{10}. The results from the simulations in the main text (x-axis) are compared to simulations with an order of magnitude smaller minimum biomass for both skewed (left) and broad (right) shaped reaction norms. The black line is the 1-1 line. Filled in data points represent increasing ΔSSTs and open data points are decreasing ΔSSTs. The minimum biomass impact is significant at the 95% CI with an average increase in offset from Q_{10}, for both reaction norm shapes.

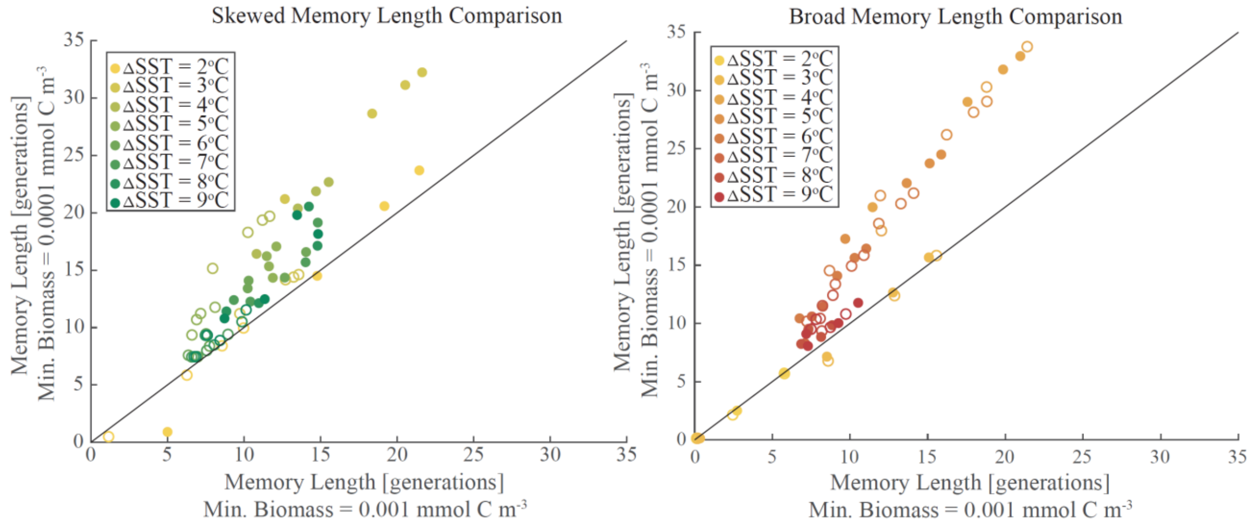

Figure S6. The impact of minimum biomass on memory length. The results from the simulations in the main text (x-axis) are compared to simulations with an order of magnitude smaller minimum biomass for both skewed (left) and broad (right) shaped reaction norms. The black line is the 1-1 line. Filled in data points represent increasing ΔSSTs and open data points are decreasing ΔSSTs. The minimum biomass impact is significant at the 95% CI with an average increase in memory length for both reaction norm shapes. However, the pattern of moderate ΔSSTs exhibiting the longest memory effects were robust across all simulations.

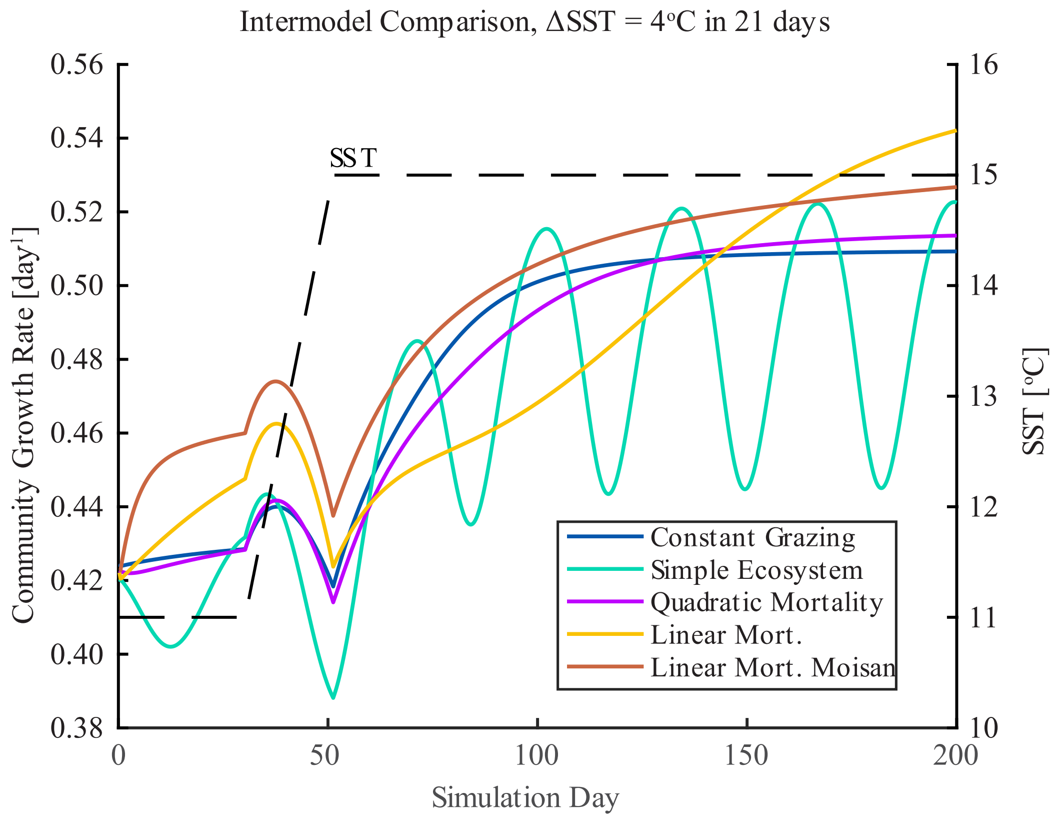

Figure S7. Comparison between different ecosystem model results for community growth for an idealized simulation with an increase of 4 °C over 21 days. For community growth rates, all models show similar qualitative results indicating a decrease in growth rate over the transient conditions culminating in a growth rate minimum when SSTs stabilize.

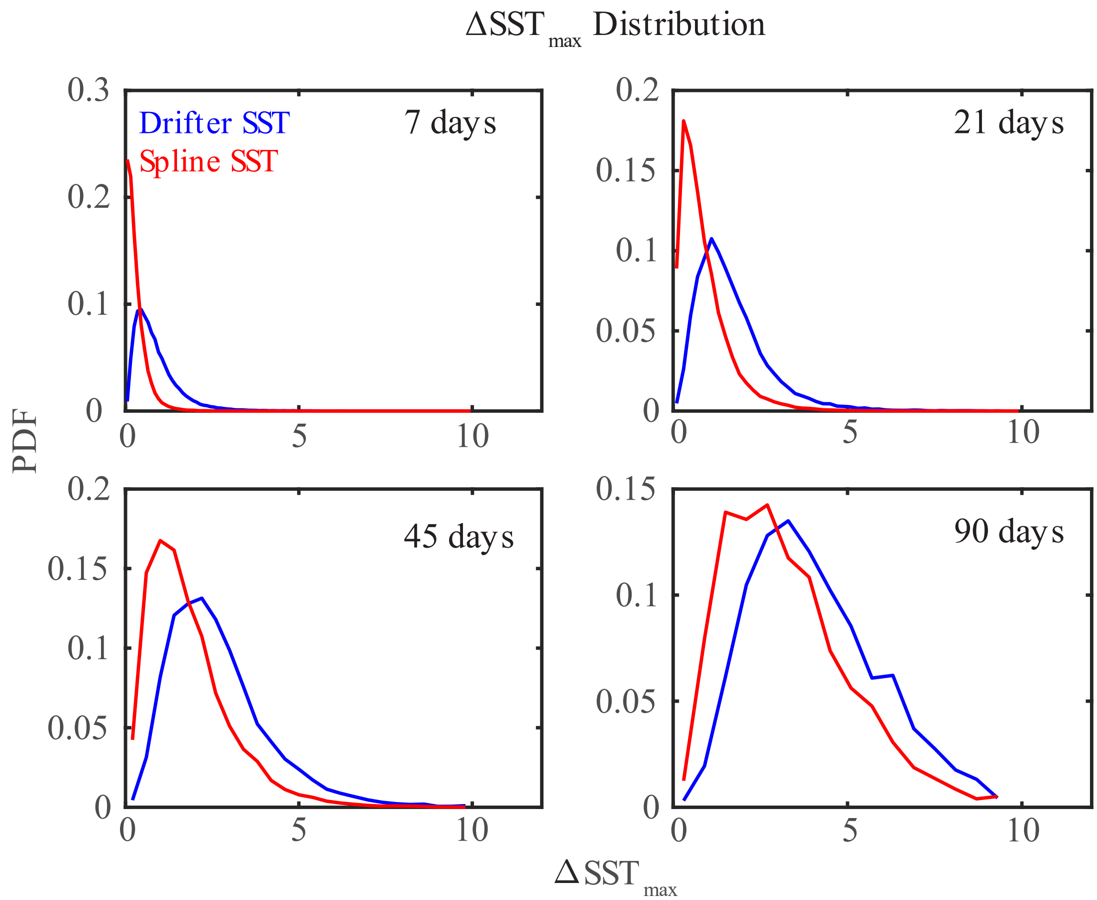

Figure S8. Probability density functions of the absolute value of the maximum change in SST over 7, 21, 45, and 90 days for the drifter trajectories (blue) and the smoothed splines of the trajectory SSTs (red).

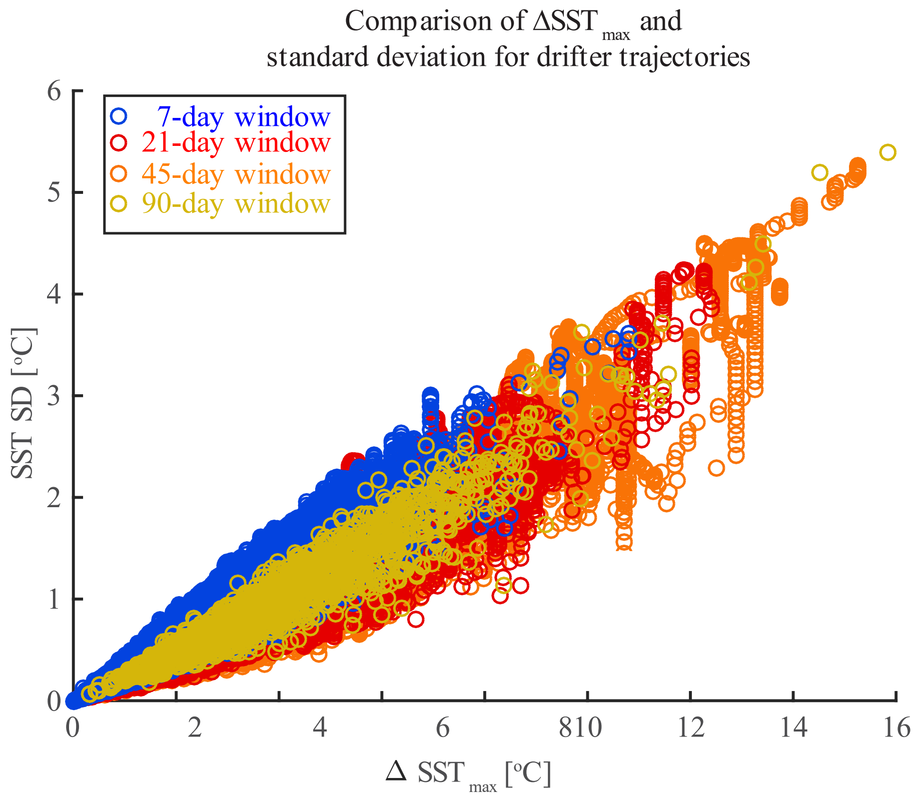

Figure S9. The standard deviation (1σ) as a function of ΔSST_{max} over different Δt_{max} windows. ΔSST_{max} drives the variability across the Δt_{max} window lengths.

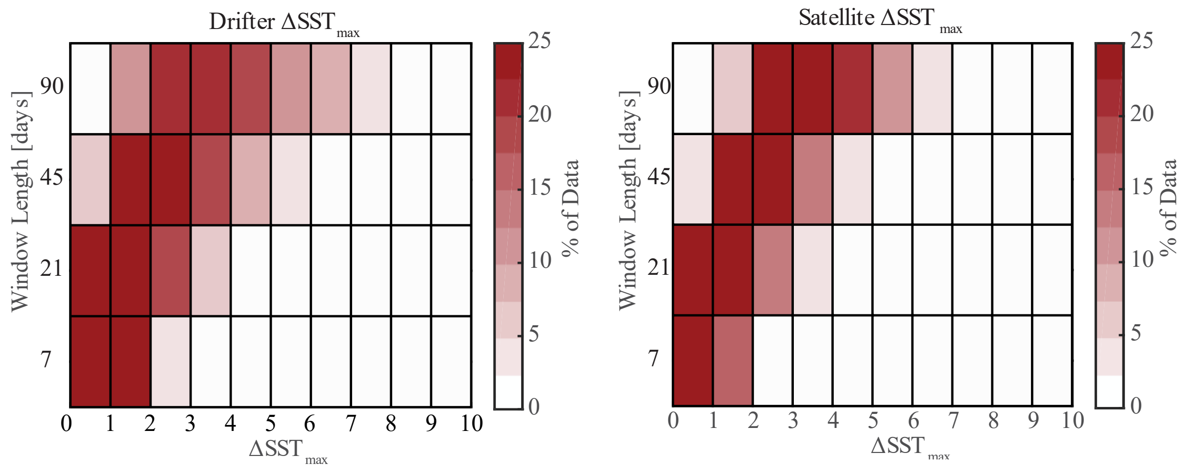

Figure S10. Results from SST variability analyses for the drifter (left) and satellite trajectories (right) showing most common SST changes for each time window. Data are presented as total percent of data that fall within that ΔSST bin for the window length. Each row sums to 100%. Although the magnitudes of variability are similar, the nature of that variability is different with the Lagrangian reference frame experiencing more variability consistent with longer memory effects.

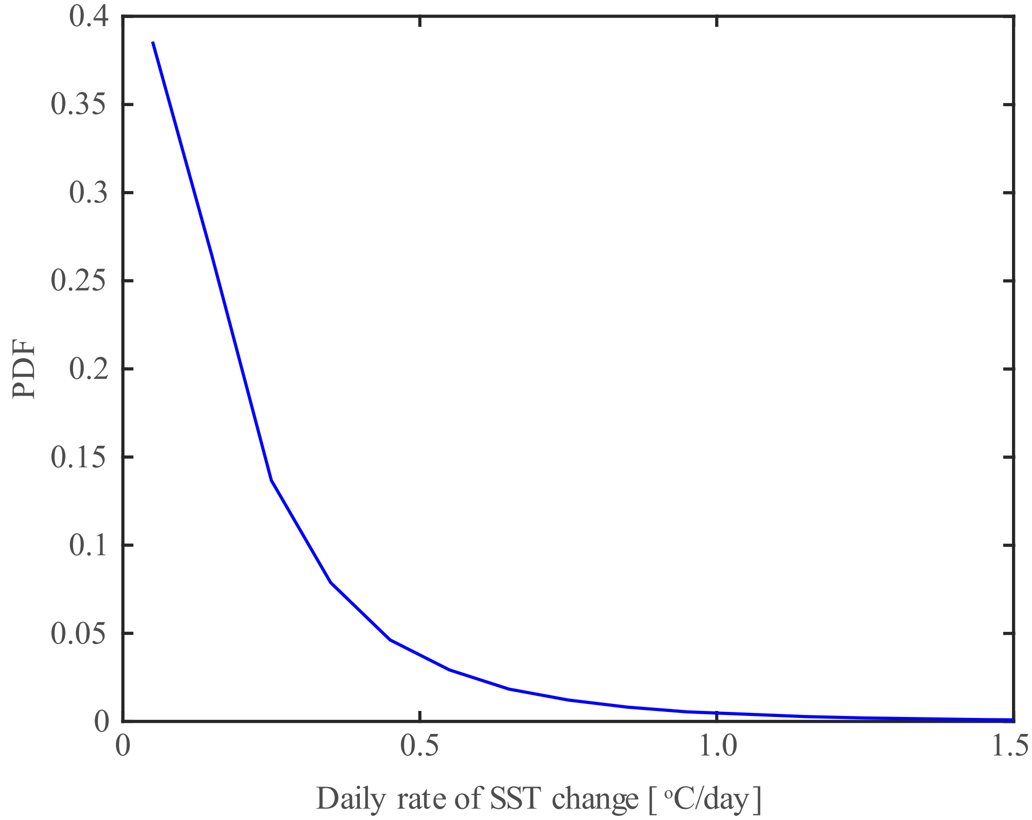

Figure S11. Daily rates of SST change for drifter trajectories. The rates of change were calculated as the range of the recorded SST values over a 1-day moving window for a total of n = 781,749 data points for 197,100 days.

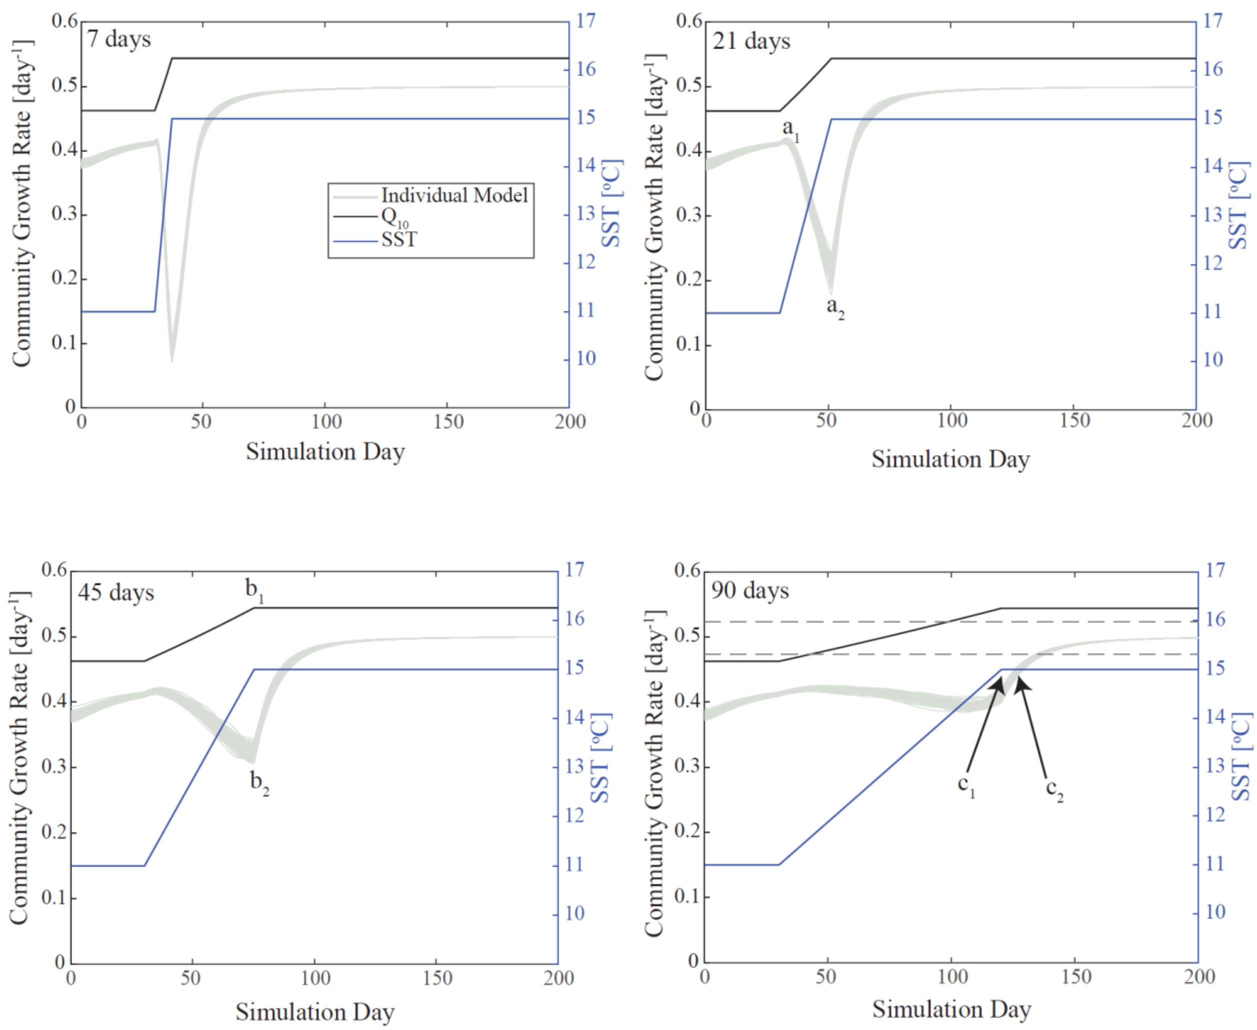

Figure S12. Community growth rates for each of the 100 simulations (grey lines) for an increase of 4°C over 7, 21, 45, and 90 days for skewed shaped reaction norms. The black line is the Q_{10} simulated community growth rate and the blue line is the SST profile for the simulation. The locations marked a_1 and a_2 in the 21-day panel represent the timesteps used to calculate the percent change in growth rate associated with transient SSTs as shown in Figures 2b. This metric was calculated as $(a_1 – a_2)*100/a_1$. The locations marked b_1 and b_2 in the 45-day panel represent the timesteps used to calculate the percent difference in growth rates between the Q_{10} parameterized growth and the phenotype model as shown in Figure 2c, S11. This metric was calculated as $(b_1 – b_2)*100/b_1$. The locations marked c_1 and c_2 in the 90-day panel point to the timesteps used to calculate the memory length. The dashed grey lines represent ±5% of the final, stabilized community growth rate which was used as the threshold for the memory effect which was defined as the time in days between c_1 when SSTs stabilize and c_2 when the community growth rate crosses the threshold.

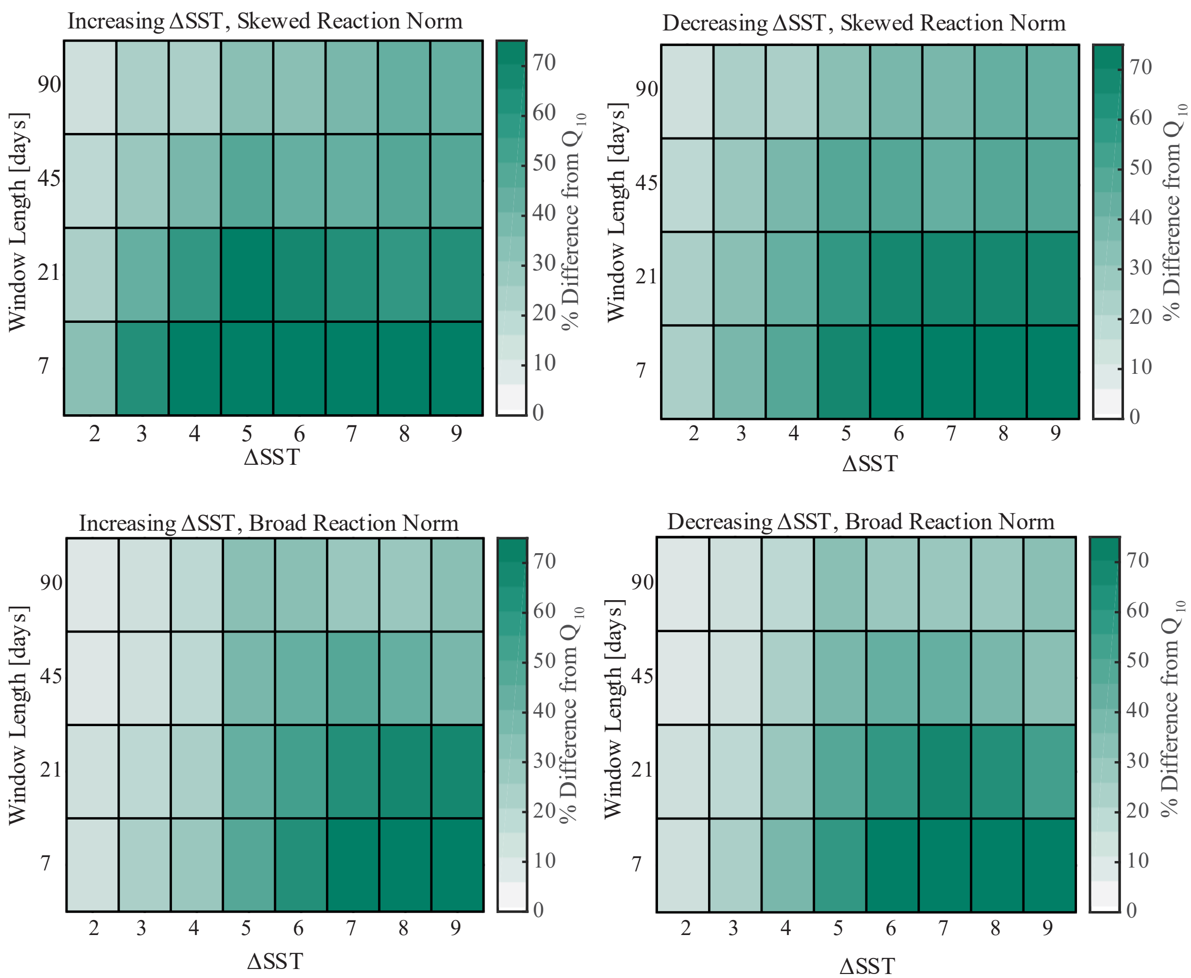

Figure S13. Full results for the percent difference from Q_{10} at the timestep when SSTs stabilize in the idealized simulations for the skewed shaped reaction norms (top row) and the broad shaped reaction norms (bottom row) under both increasing SSTs (left column) and decreasing SSTs (right column).

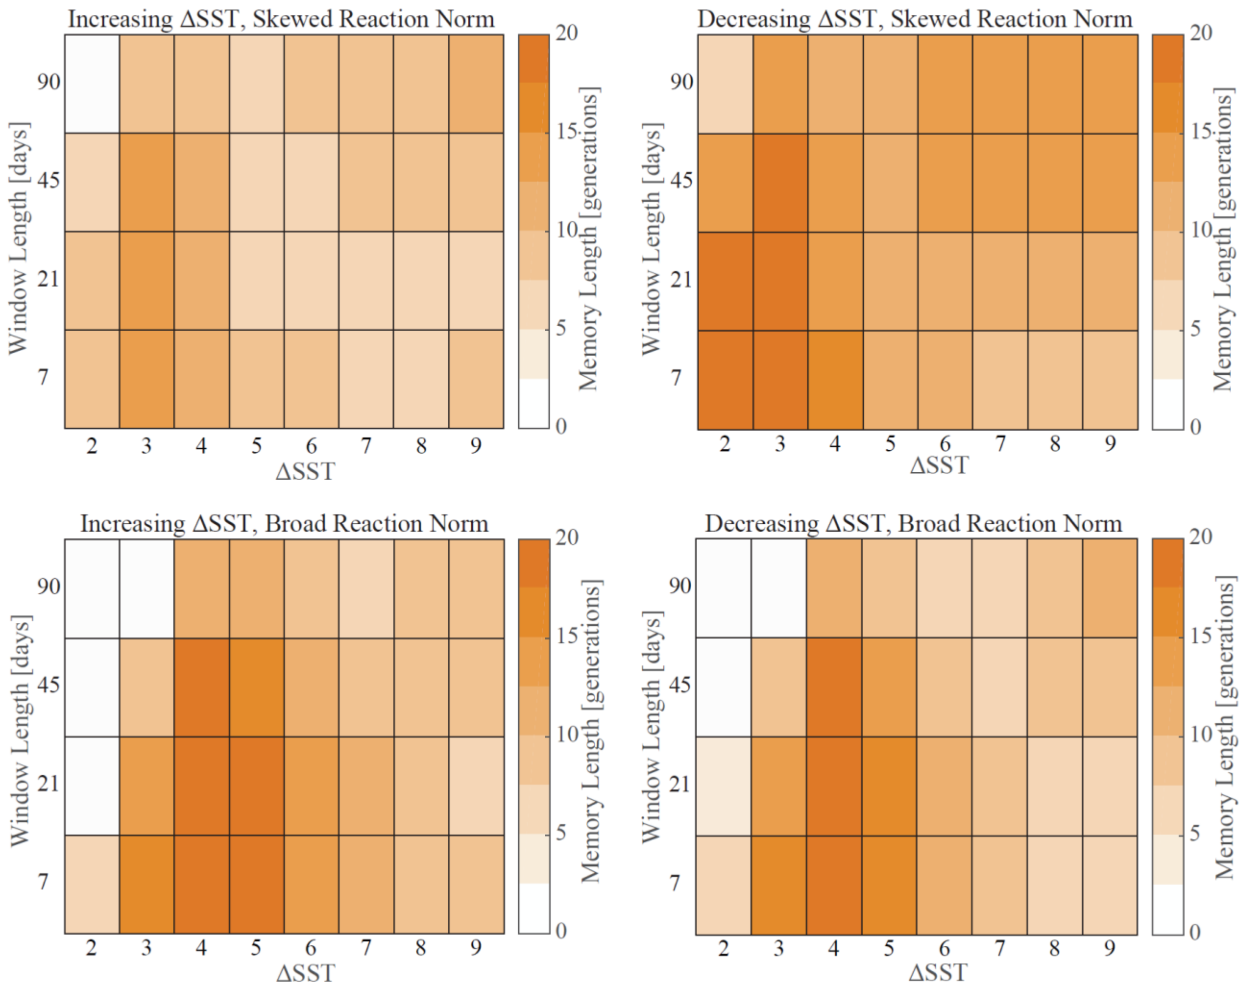

Figure S14 Full results for the length of the memory effect in the idealized simulations for the skewed shaped reaction norms (top row) and the broad shaped reaction norms (bottom row) under both increasing SSTs (left column) and decreasing SSTs (right column).

.

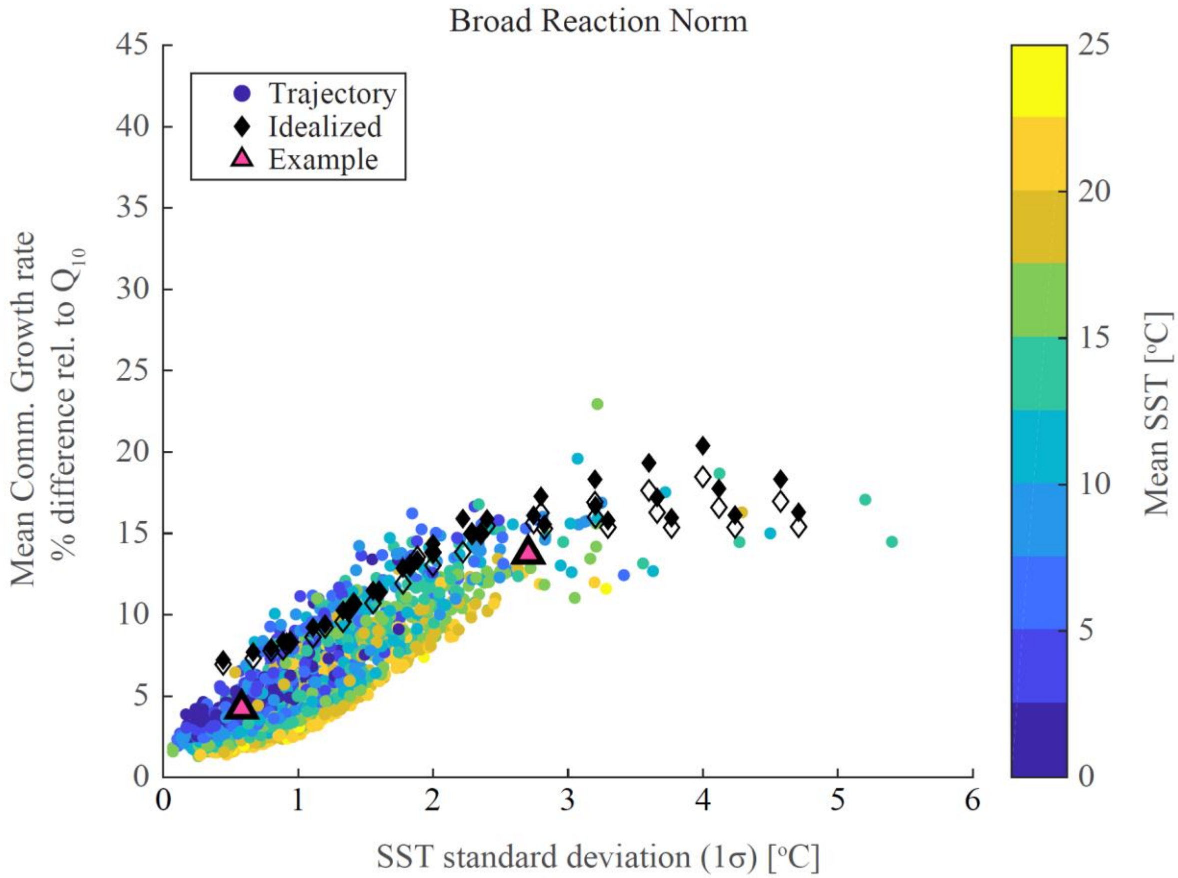

Figure S15. The 90-day average percent difference between community growth rates determined via the Q_{10} method and the phenotype-based model versus standard deviation (1σ) of SST over the 90-day trajectory. Drifter data are represented by circles colored according to their mean SST. Black diamonds represent the first 90 days of the idealized trajectories; filled diamonds are the idealized trajectories for which SSTs increase and open black diamonds are idealized trajectories with decreasing SSTs. Pink triangles represent the two example trajectories from Figure 1 in the main text.

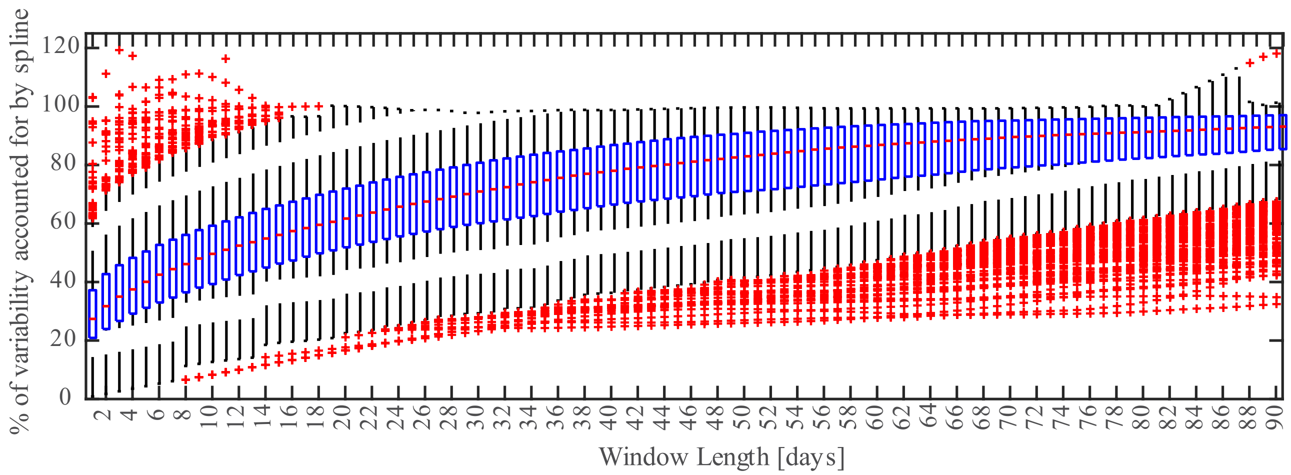

Figure S16. Box plots of the percent of the SST variability in the drifter trajectory that is accounted for by the smoothed spline. Each of the 2,190 90-day drifter and spline trajectories was broken up into windows in 1-day increments from 1 to 90 days. The standard deviation of the drifter trajectory is the sum of the standard deviation of the smoothed spline plus some noise term. From this, the variability accounted for by the spline for each window, for each trajectory was recorded with the results shown. As expected, over longer window lengths the spline accounts for higher percentage of the overall variability.

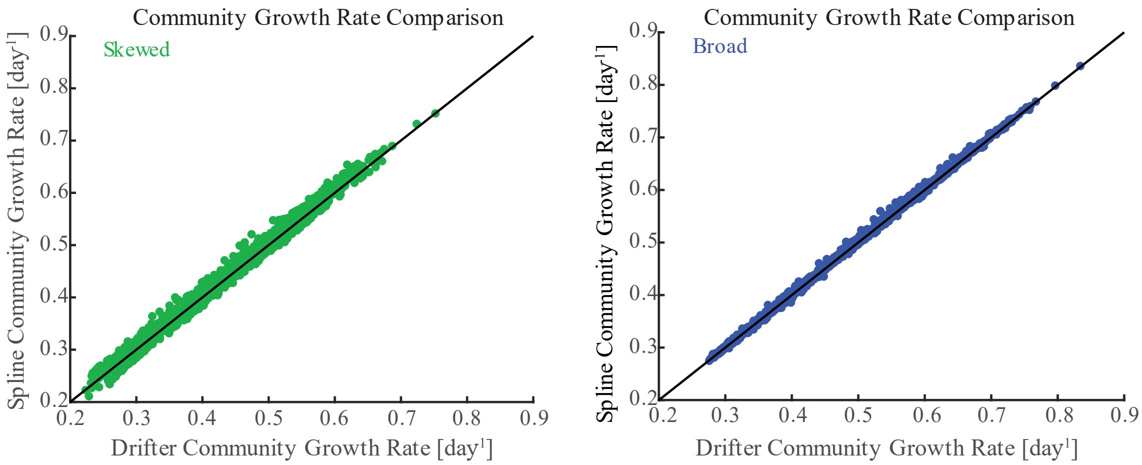

Figure S17. Comparison of mean community growth rate over the entire 90-day trajectory for the real trajectories and their spline simulations for skewed (left) and broad (right) shaped reaction norms. With each reaction norm shape, smoothing the small-scale noise did not impact overall biomass-weighted community growth rates (95% CI, t-test) further supporting that small-scale noise does not induce a memory effect.

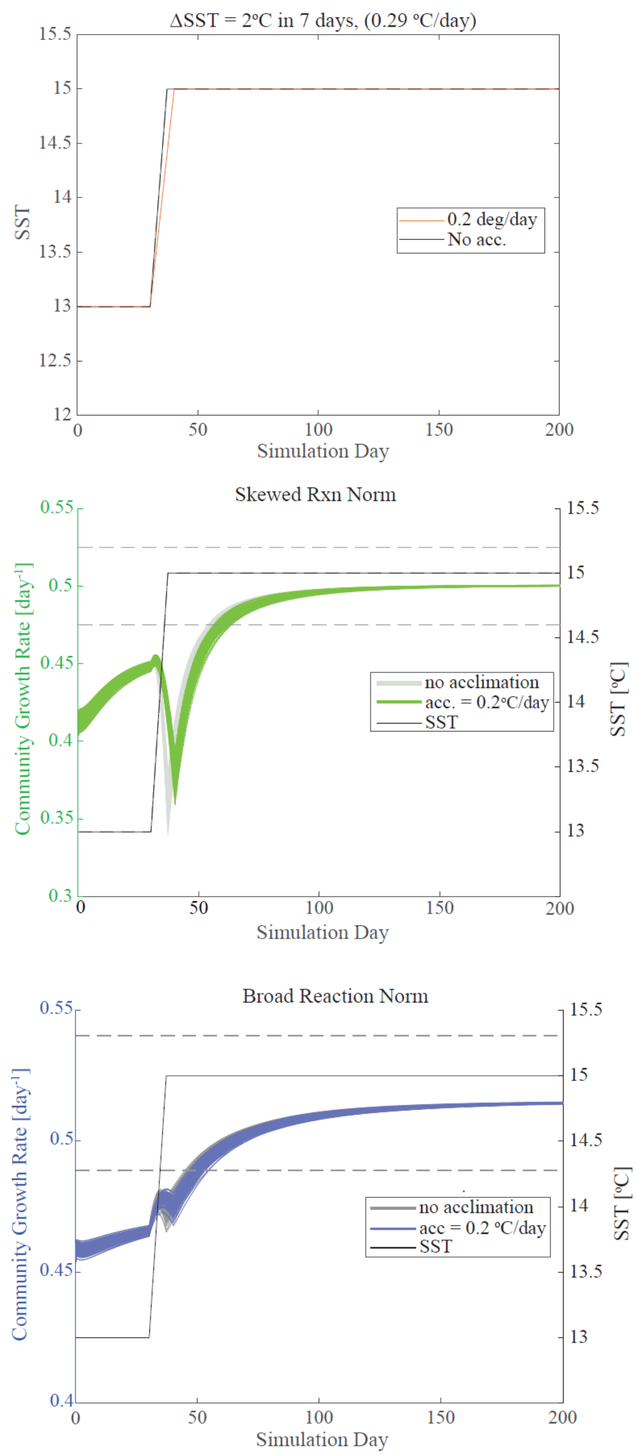

Figure S18. Example of impact of acclimation. (Top) Example idealized SST trajectory of changing 2 °C in 7 days with acclimation rates of 0.2 °C day^{-1} and 0.3°C day^{-1}. Other acclimation rates not shown as they plot along the No Acclimation line because those rates are faster than the rate of change. (Middle) Community growth rates for skewed reaction norms over each of the simulations for the no acclimation simulations (grey lines) and the simulations with an acclimation rate of 0.2°C day^{-1} (green lines). Dashed lines represent the thresholds used to calculate the memory length. (Bottom) Same as the middle panel but for broad reaction norms.

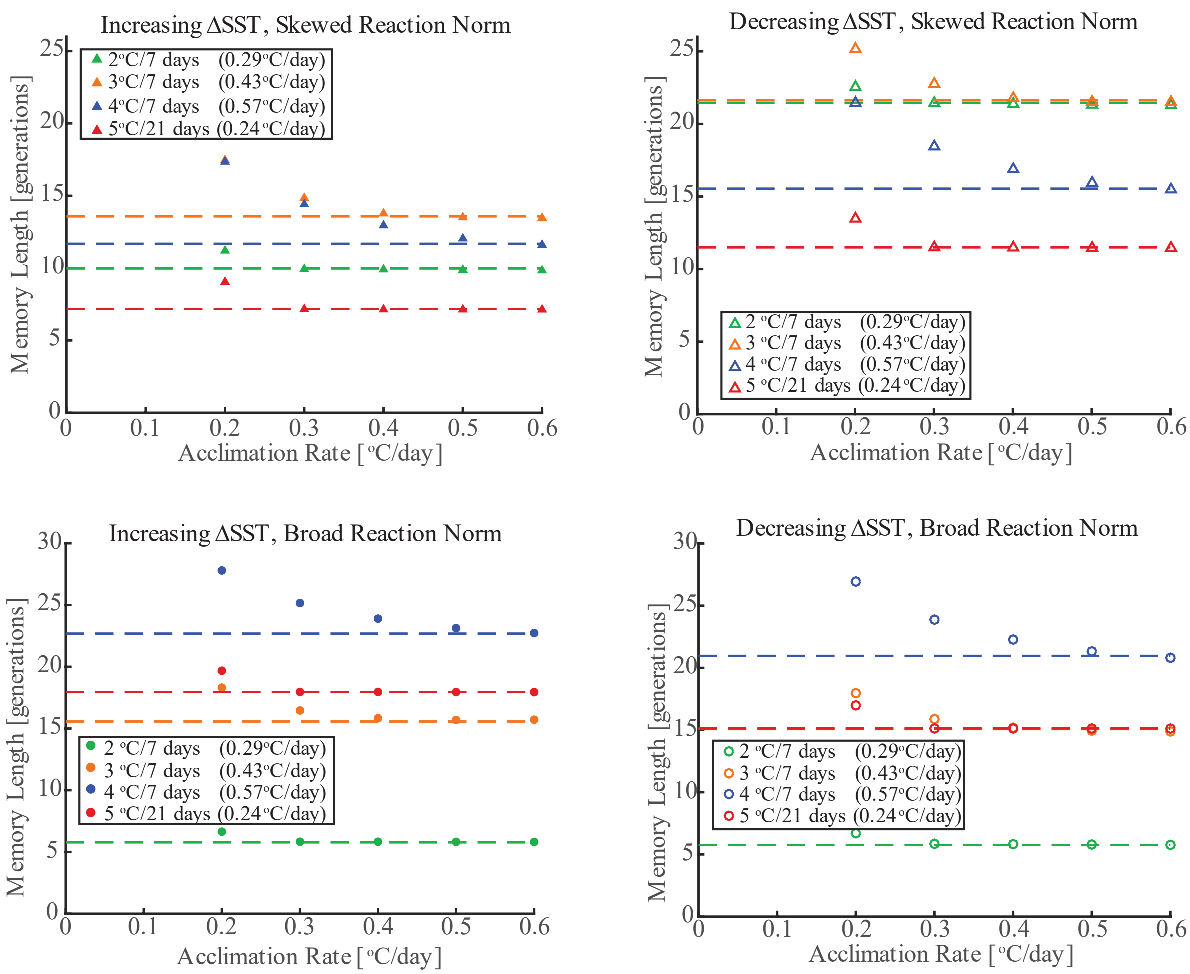

Figure S19. Impact of acclimation on memory length on the skewed reaction norms (top row) and the broad reaction norms (bottom row) in both the increasing (left column) and decreasing (right column) ΔSST directions. Dashed lines represent the memory lengths calculated for the simulations that did not incorporate acclimation. When acclimation rates are greater than or equal to the rate of SST change, there is no difference between the simulations that incorporated acclimation and those that did not.

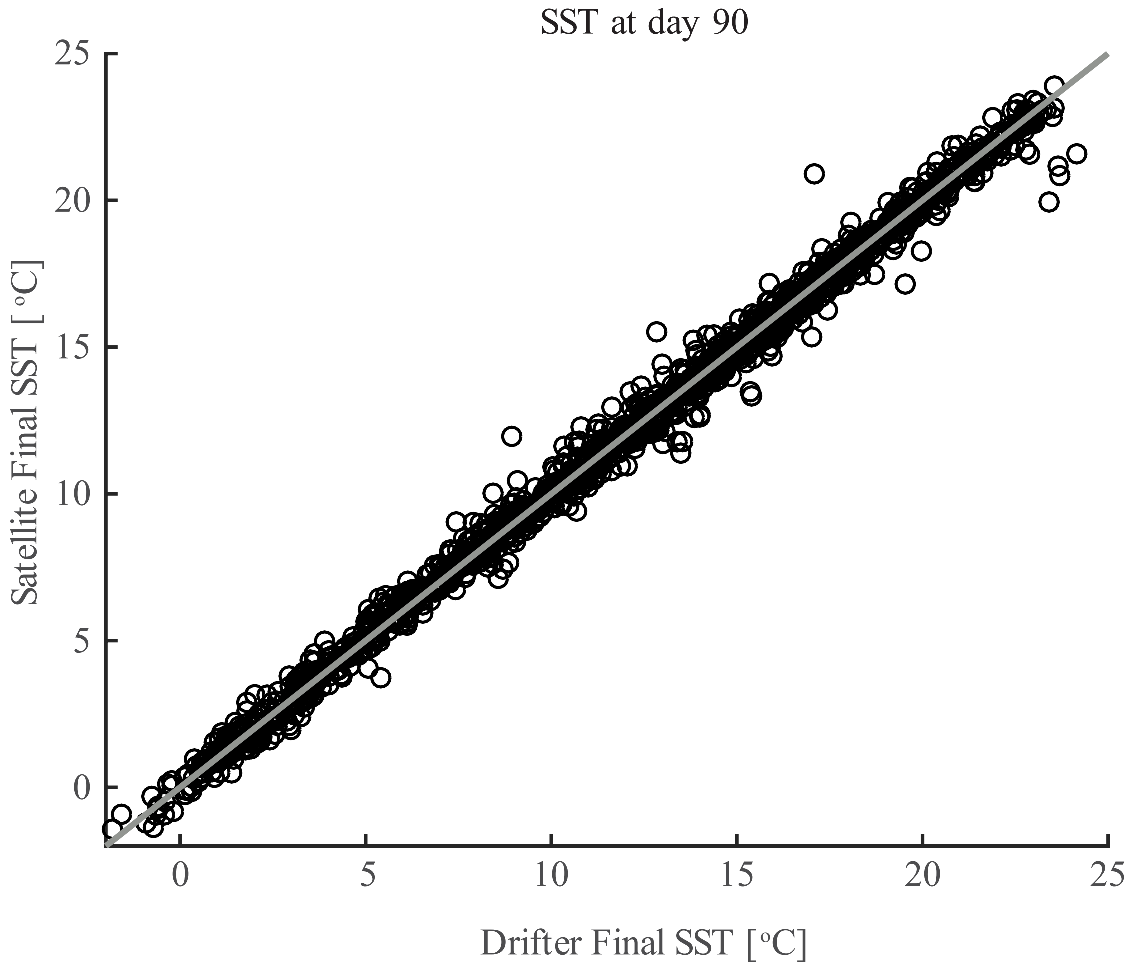

Figure S20. Comparison of the final SST for the drifter and the satellite data. The data from both sources represent the same location in space and time so the data should be similar and in fact, are not statistically different from one another (ttest, 95% CI). The grey line represents the 1-1 line.

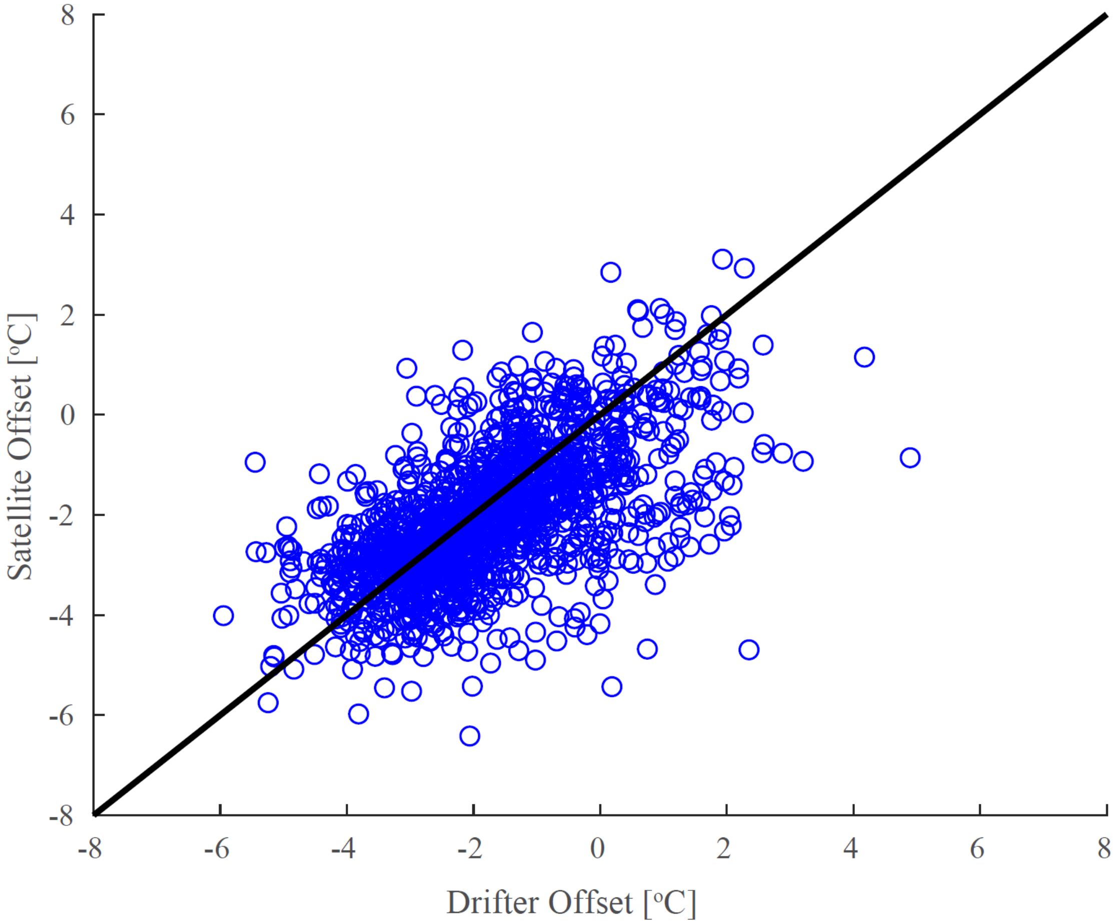

Figure S21. Same as Figure 4 in the main text but for broad shaped reaction norms. The impact of Lagrangian and Eulerian variability on community composition. Here we plot the difference between the Topt of the most abundant phenotype at the end of each 90-day trajectory and the final SST for the drifter trajectory (x-axis) and the satellite data (y-axis). The final SSTs for the drifter and satellite data are statistically identical (t-test, 95% CI). Therefore, deviations from the 1:1 line demonstrate the impact of a Lagrangian versus Eulerian reference frame on community composition.

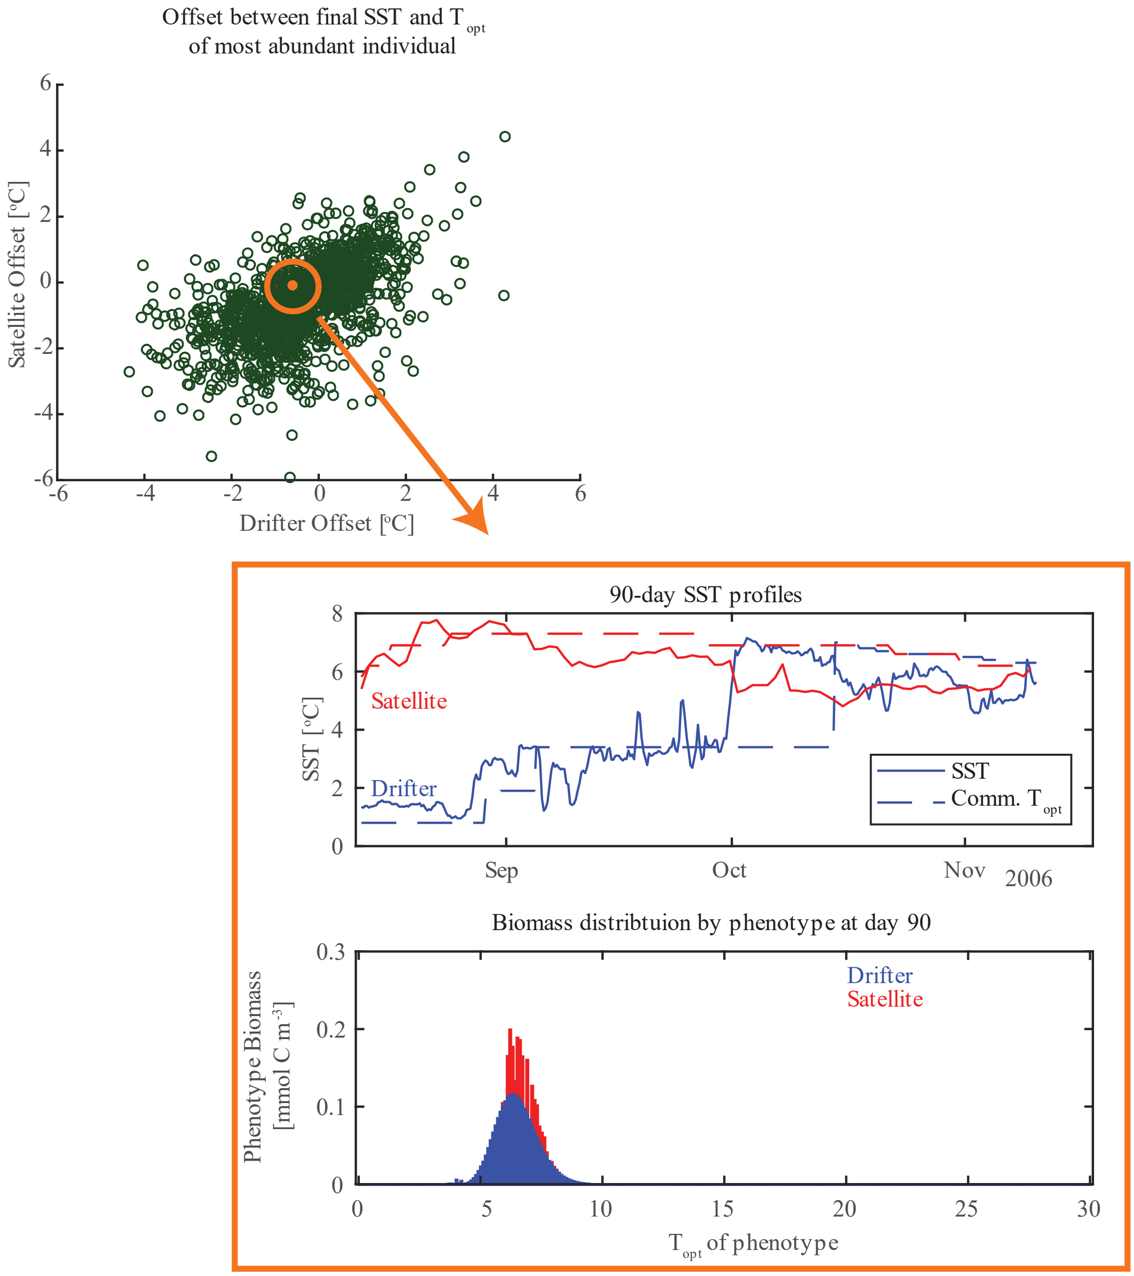

Figure S22. The impact of SST variability on community composition. (Top) An example 90-day drifter trajectory and the satellite SST data for the final location of the drifter over the same 90 days shown as solid lines. The dashed lines are the T_{opt} of the most abundant phenotype at each timestep. (Bottom) The biomass of each phenotype with a skewed shaped reaction norms at day 90 for the satellite and drifter trajectories. In this example, the offset between the final SST is -0.60°C for the drifter and -0.09°C for the satellite data. The difference in the magnitude of the offset between the two data sets represents the difference in the variability of the SSTs. However, in this example, the satellite SSTs stay relatively constant whereas the drifter SSTs experience a rapid increase of 3.5°C in 4 days beginning Sept. 29. Because the drifter SSTs remain relatively constant through the end of the 90 days, the community is able to adjust to the new environment before the end of the simulation which results a community T_{opt} that reflects the SSTs at day 90 for both the satellite and the drifters.

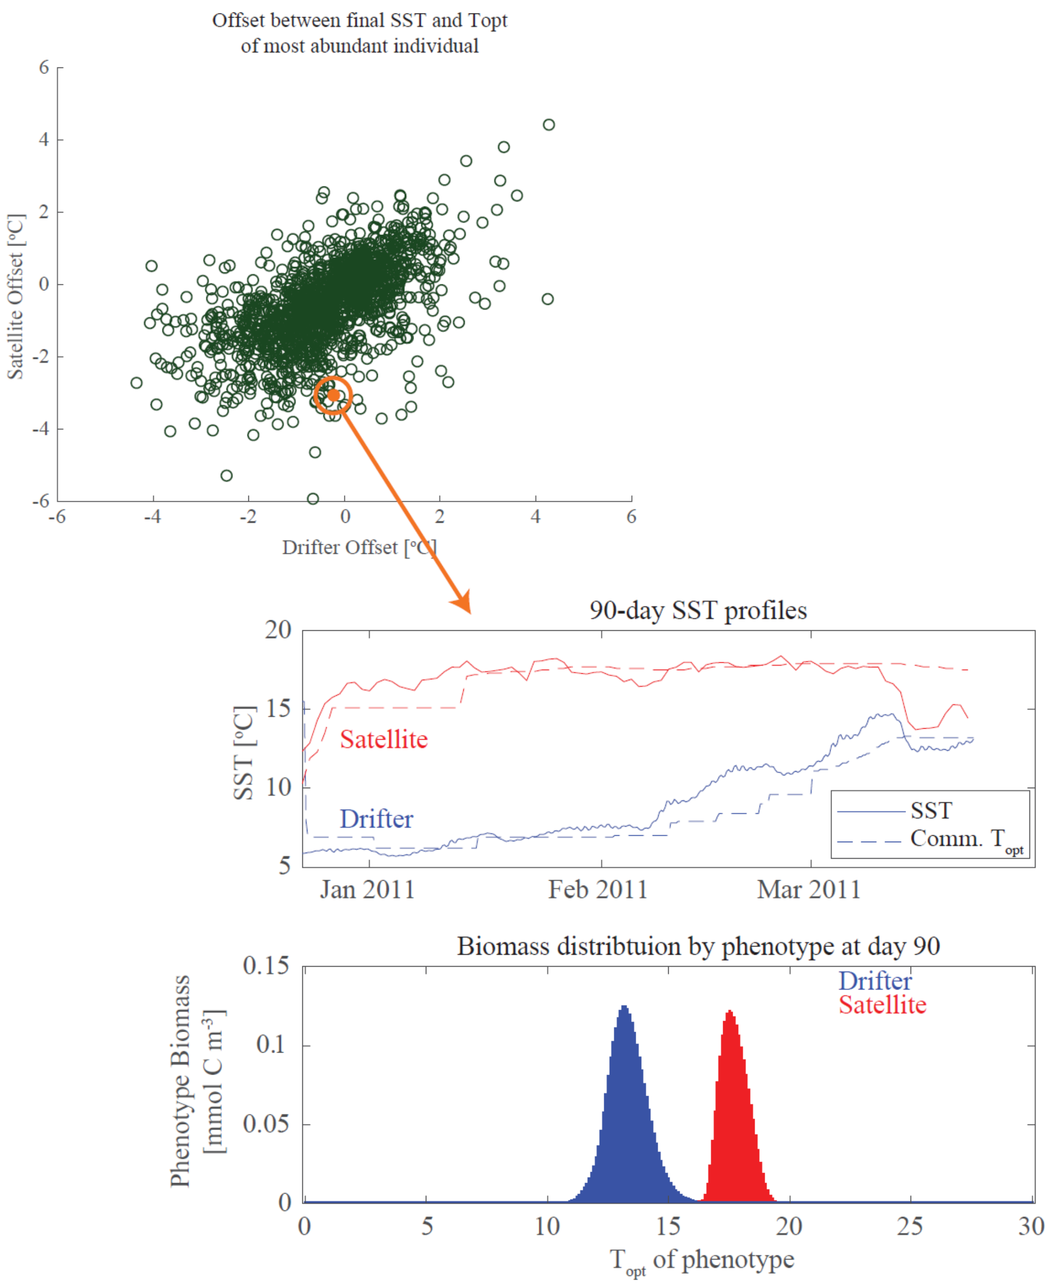

Figure S23. The impact of SST variability on community composition. (Top) An example 90-day drifter trajectory and the satellite SST data for the final location of the drifter over the same 90 days shown as solid lines. The dashed lines are the T_{opt} of the most abundant phenotype at each timestep. (Bottom) The biomass of each phenotype with a skewed shaped reaction norms at day 90 for the satellite and drifter trajectories. In this example, the offset between the final SST is -0.23ºC for the drifter and -3.1ºC for the satellite data. The difference in the magnitude of the offset between the two data sets represents the difference in the variability of the SSTs. Here, the drifter SSTs gradually increase over the 90 days which allows the community to continuously track the changes in SST whereas the satellite SSTs are relatively stable and then rapidly decrease from 17.7ºC on March 10 to 13.8ºC on March 17. Due to the long memory effect associated with this rate and magnitude of change, the community was not able to track the SST change which resulted in a large offset between the final SST and the T_{opt} of the most abundant phenotype at day 90.