Resource-explicit interactions in spatial population models

Supporting Information

RESOURCE NODE PLACEMENT METHODS

Three resource node placement methods were used in the models assessed in this manuscript. These methods were: placing resource nodes at the center of each square in a uniform square tiling of the landscape (Fig. 2), placing resource nodes at the center of each hexagon in a uniform hexagonal tiling of the landscape (Supplemental Fig. 1), and placing the resource nodes randomly across the landscape, re-randomizing the placement of the nodes during each tick of the model (Supplemental Fig. 2). Specifically, random placement was accomplished by assigning each node an x and y coordinate randomly drawn from a uniform distribution with a minimum of 0 and a maximum of the side length of the model. For each of these placement methods, several resource node densities were tested.

In terms of code complexity, the randomized placement method is the simplest. This method allows resource nodes to be added to the model with just a few lines of code. The code to perform square and hexagonal tilings in our models is significantly more complex due to the desire for the tiling to cover any square area of arbitrary dimensions. To accomplish this, the grid of nodes is centered within the modeled area, and nodes with areas that fall partially outside the landscape are parameterized with a proportionately decreased amount of resources. For example, if tiling a 2.8 by 2.8 landscape with square nodes that are each 1 by 1, a 3 by 3 grid of nodes would be used, with the grid of nodes centered within the area, and with nodes along the edges of the area are parameterized with an appropriately reduced amount of resources given the portion of the node's area that falls within the modeled area. In the hexagonally tiled models, the code required to accomplish this task is somewhat more complex than in the square-tiled models. That said, this complexity was included in the interest of matching the direct-interaction models in this manuscript as closely as possible. In models designed from the ground up as resource-explicit that do not need to match a prior model, the landscape could be defined as having a boundary exactly matching a desired grid of nodes, removing the need for this complexity. Minor adjustment to these methods would also be appropriate if modeling a periodic area (such as a toroidal landscape), in order to ensure appropriate resource availability along the "seams" of the model.

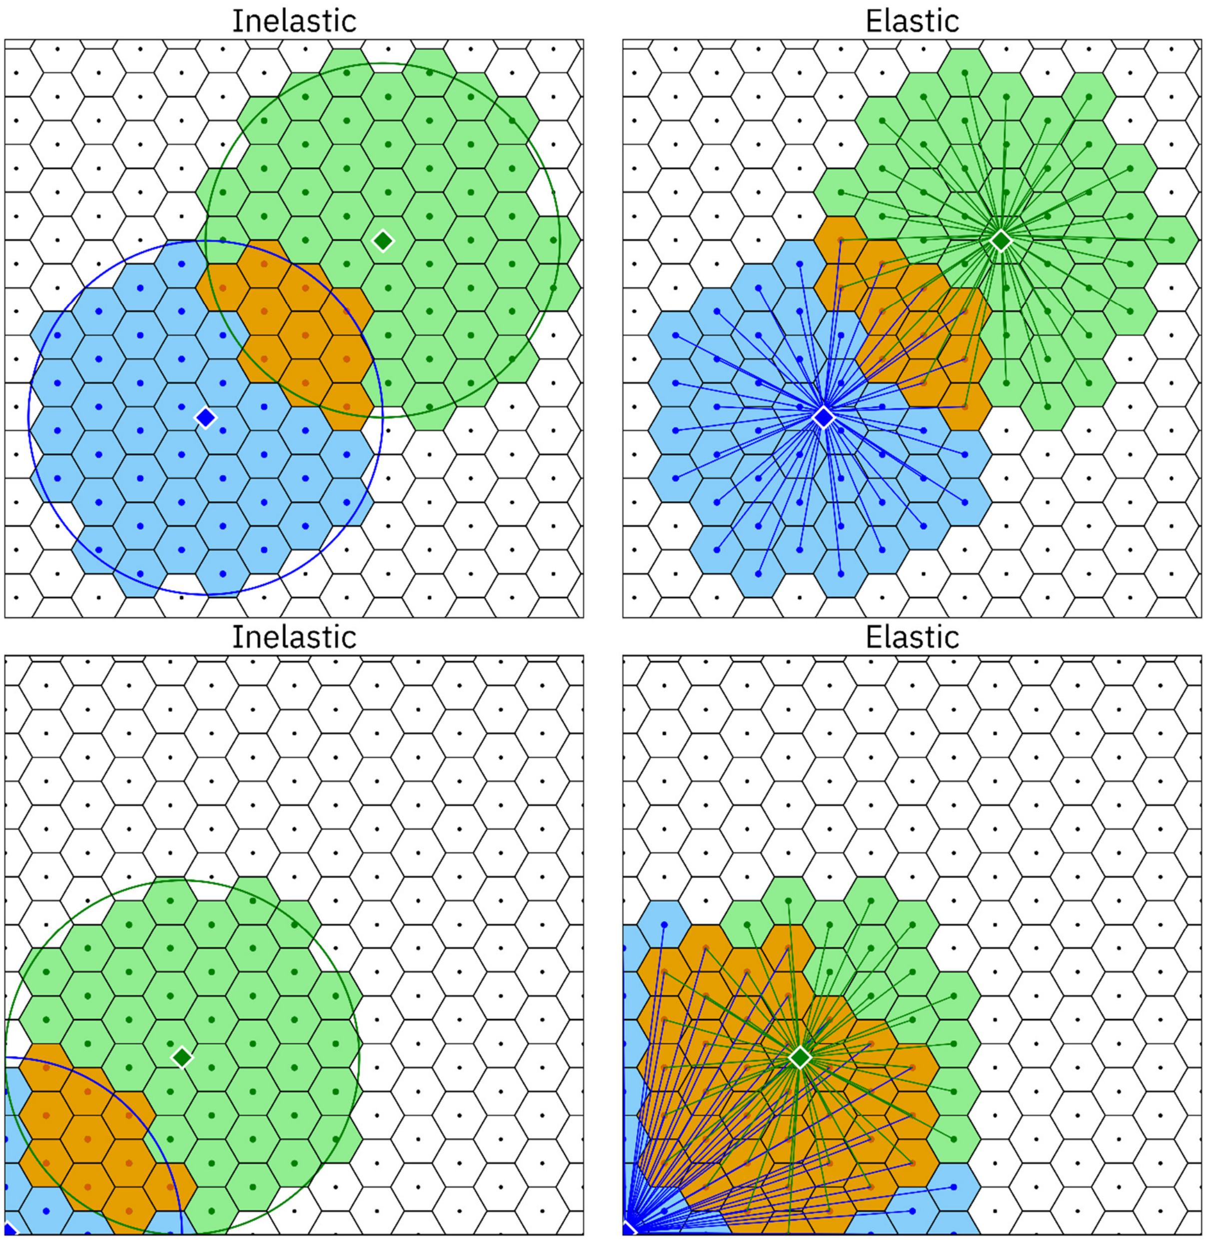

Supplemental Figure 1. Visualization of resource-explicit interaction algorithms, hexagonal tiling. Competition between two individuals (blue and green diamonds with white outlines) is determined by the portion of their foraging area that overlaps and by the interaction algorithm used in the model. The foraging areas are represented by blue and green shading; the overlapping area is shaded orange. In the inelastic model (left), individuals forage from resource nodes (dots at the center of each hexagon) within their foraging radius. In the elastic model (right), individuals forage from as many nodes as necessary to maintain a nominally sized foraging area (in this case, that area comprises 50 nodes). In upper left panel, the blue individual happens to forage from only 49 nodes, and competition between these two individuals is slightly reduced. In the corner of the landscape, the difference between the two models is much greater: the blue individual has a much smaller foraging area in the inelastic model, while the blue individual in the elastic model forages from much further away to maintain a full-sized foraging area, resulting in greater competition.

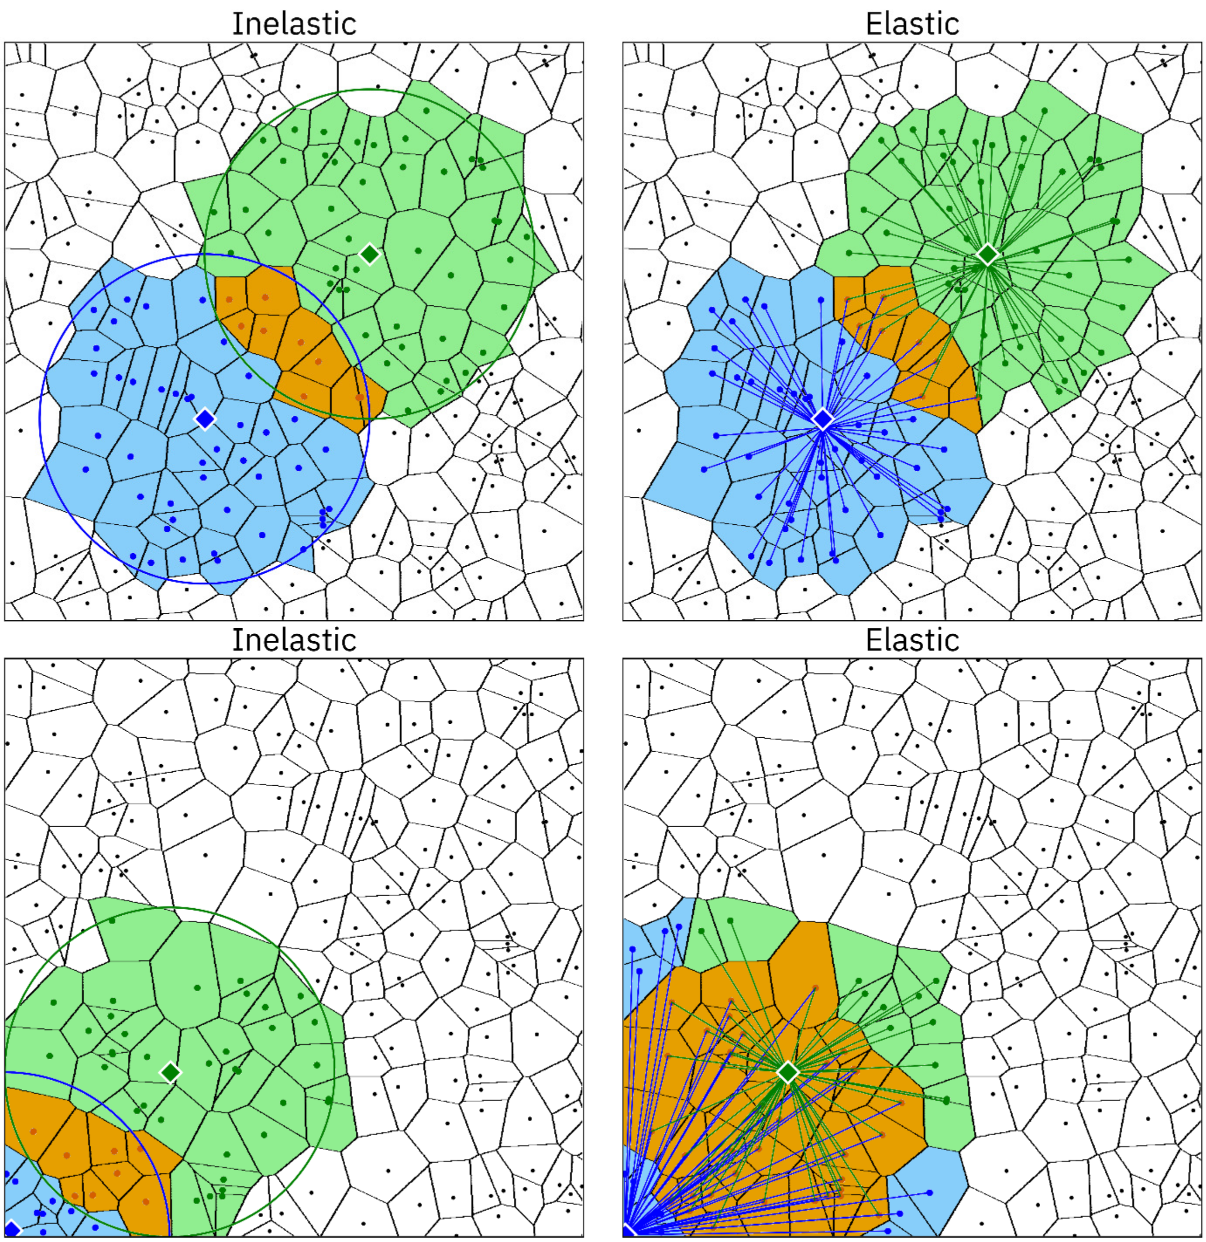

Supplemental Figure 2. Visualization of resource-explicit interaction algorithms, random node placement. Competition between two individuals (blue and green diamonds with white outlines) is determined by the portion of their foraging area that overlaps and by the interaction algorithm used in the model. The foraging areas are represented by blue and green shading; the overlapping area is shaded orange. In the inelastic model (left), individuals forage from resource nodes (dots at the center of each polygon) within their foraging radius. In the elastic model (right), individuals forage from as many nodes as necessary to maintain a nominally sized foraging area (in this case, that area comprises 50 nodes). When nodes are randomly distributed, the difference between the inelastic and elastic method even in the interior areas is greater. For example, the blue individual in the upper left panel forages from 53 nodes, and the green individual in the bottom left panel forages from 60 nodes, whereas in the elastic model, all individuals forage from 50 nodes, regardless of their position.

Choice of node density and placement method. As demonstrated in the Results section, the choice of resource node density effects the runtime of the model as well as the magnitude of variance in interaction strengths within the model. An increased node density tends to yield a tighter distribution of interaction strengths, but runs more slowly. The choice of node density must therefore be made based on what kind of interaction strength variance can be can be considered acceptable and how fast the model needs to run.

The choice between a regularly tiled model or a model with randomly distributed nodes is also fairly straightforward. Models with regular tilings have tighter distributions of interaction strengths, and run slightly faster compared to models where the node positions are re-randomized during each tick of the model. However, a model with randomized resource positions that are not re-randomized could be used to represent a random heterogeneous landscape that is uniquely generated each time the model is run. Random node placement could also be combined with a landscape map to allow for the simulation of a specified heterogeneous landscape without any pre-calculation step. Additionally, more complex node placement strategies may provide desirable ways to represent realistic resource variability. For example, nodes might be placed according to a Poisson-Disc Sampling, in which entities are randomly placed, but are not placed closer to one another than a specified minimum distance.

When minimizing variance in interaction strength is a priority, a regular tiling of resource nodes is the best option (Supplemental Fig. 3; compare Figs 3 and 4 versus Fig. 5). This holds true both for simulations that are intended to emulate a homogenous space (as is common in direct-interaction models) and for simulations involving heterogeneous landscape maps (in which case pre-calculating the appropriate amount of resources at each node and using a dense tiling will result in the resource-explicit model matching the landscape as closely as possible).

The choice between a hexagonal grid and a square grid is less consequential. The two are almost interchangeable. Since a square grid is simpler to implement, it may be the right choice for most models. In most cases, the standard deviation in interaction strengths as compared to the circle-intersection function with one tiling or another at a given density varies by no more than a percent or two (Supplemental Fig. 3).

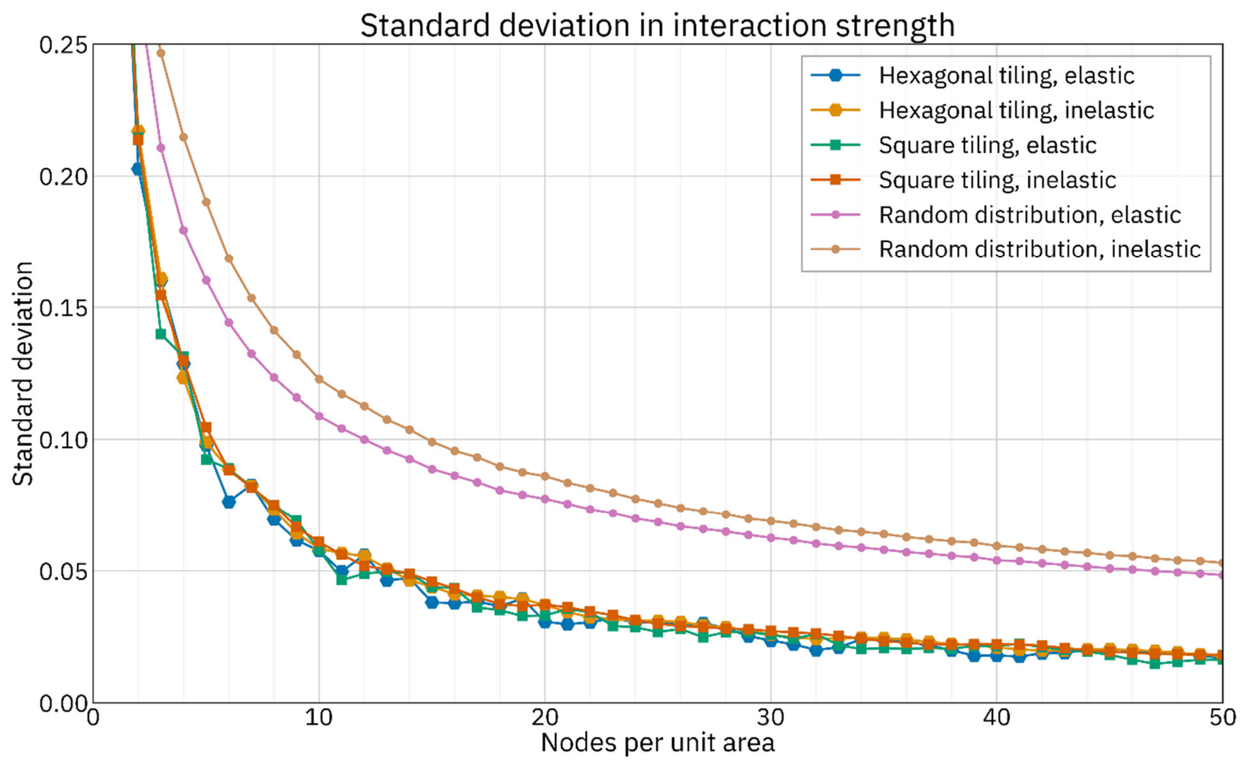

Supplemental Figure 3. Standard deviation of the differences in interaction strength between resource-explicit models and the circle-intersection function, square tiling. Two million pairwise interaction strengths measured by the circle-intersection function were subtracted from those measured in the resource-explicit models to yield a distribution for each model. The standard deviation was then measured for that distribution. This was repeated with tiling density ranging from 1 to 50 nodes per unit area. Note: while perhaps surprising, the cases where a denser tiling has a higher standard deviation are not due to measurement error (e.g., a hexagonally tiled elastic model with 7 nodes per unit area has a higher standard deviation than the same model with a density of 6 nodes per unit area).

There are only three possible ways to tile a plane with edge-to-edge congruent regular polygons: using equilateral triangles, squares, or hexagons (Birch et al., 2007; Grünbaum & Shephard, 1977). In other ecological modeling contexts, such as modeling an area by using an array of panmictic demes linked by migration, it is well accepted that spatial artifacts are smaller when using hexagonal grids compared to using square grids (Birch et al., 2007). Without going into complete detail, the reason the tiling does not matter as much in the resource-explicit method is as follows: in our method, the *area* of each shape (be it a hexagon or square) is determined by node density parameter. In linked panmictic models, the *distance* between nodes is fixed by the model (with the distance between what are considered to be different populations being related to the dispersal characteristics of the species being modeled). Given a fixed distance between nodes, a hexagonal tiling is a denser packing, with each hexagon therefore having a smaller area than squares would have in a square tiling with the same distance between nodes. Thus, a hexagonal tiling offers a large advantage in linked panmictic models, but is not that different from a square tiling in our resource-explicit models. However, in some cases, a hexagonal tiling has some desirable properties in terms of minimizing unequal availability of resource nodes in the inelastic model (see Supplemental Figs 14-19).

DYNAMICS OF THE RESOURCE-EXPLICIT MODELS

In the Results section, interaction strengths in the resource-explicit models were assessed by measuring pairwise interactions between randomly placed individuals. The interaction strengths measured in the resource-explicit models were then subtracted from the strengths measured by the circle-intersection function. This analysis was performed using a square tiling of the landscape (Fig. 4), a hexagonal tiling of the landscape (Supplemental Fig. 4), and with resource nodes randomly placed, with new positions drawn for each interaction measured (Supplemental Fig. 5).

A further analysis was performed that compared the overall interaction forces experienced by individuals in the resource-explicit models to that calculated by the circle-intersection function. This analysis was performed using a square tiling of the landscape (Fig. 5), a hexagonal tiling of the landscape (Supplemental Fig. 6), and with resource nodes randomly placed, with new positions drawn for each interaction measured (Supplemental Fig. 7).

We performed an additional assessment of interaction strengths by measuring an "interaction field" between a focal individual and a dense grid of points within interaction range. This analysis was performed with a focal individual placed at a coordinate that minimized the variance between the resource-explicit model and the circle-intersection function, and again at a coordinate that maximized the variance. This analysis was conducted with a hexagonally tiled landscape as well as a square-tiled landscape, each at a density of 12, 25, and 50 nodes per unit area, each with both the inelastic and elastic method (Supplemental Figs 8-13).

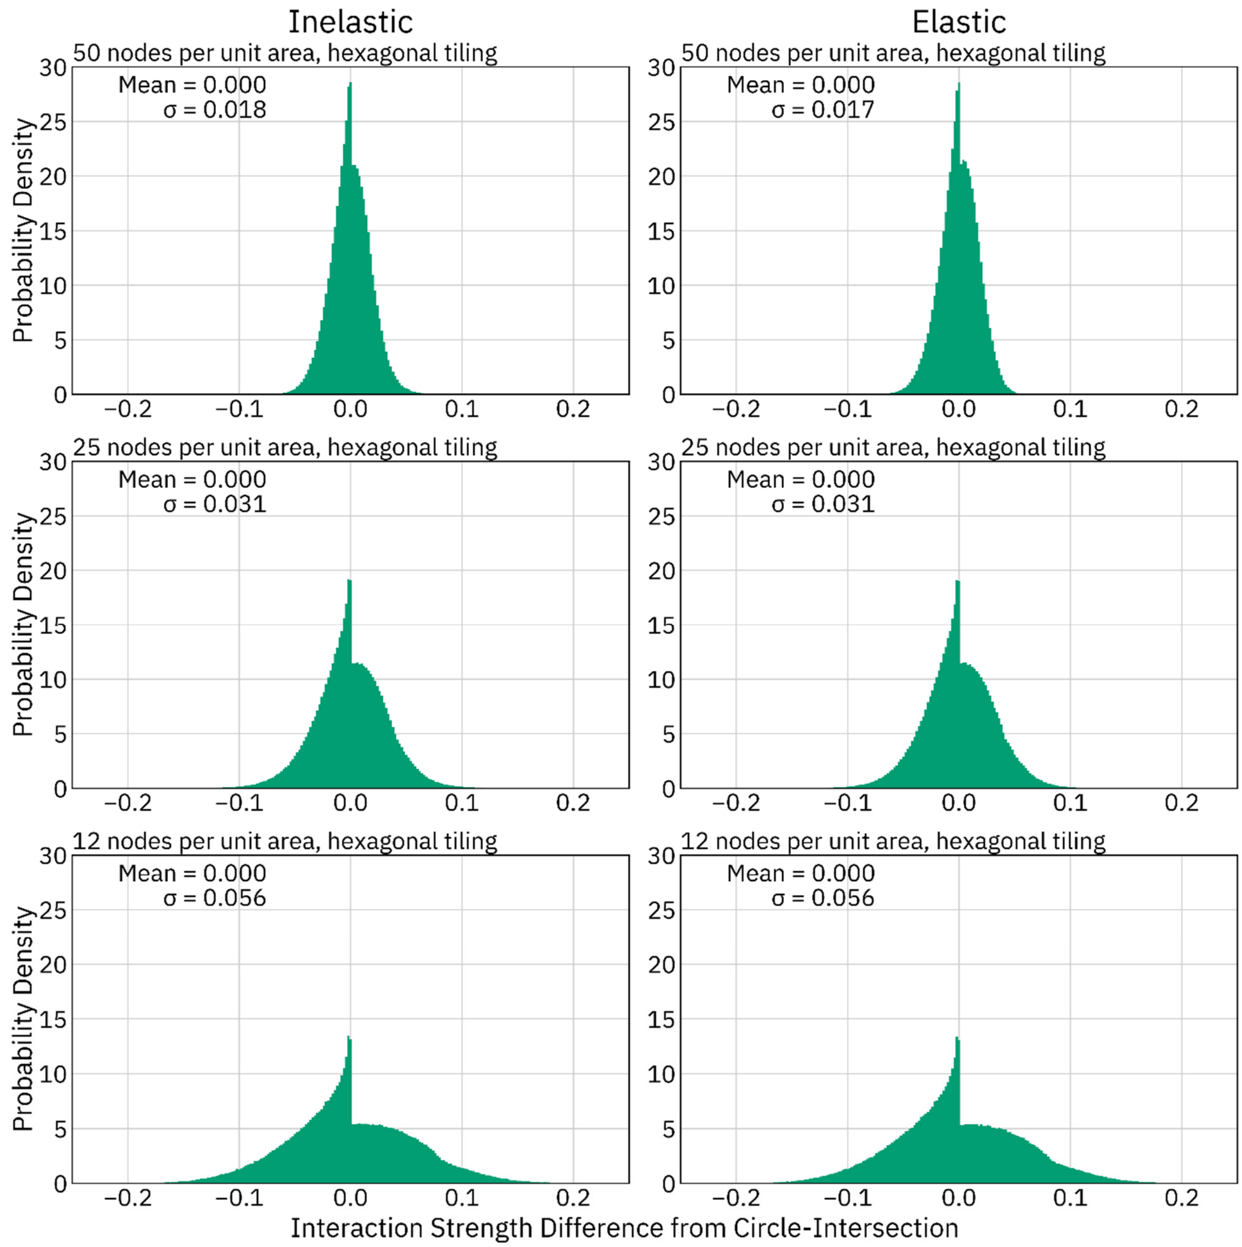

Supplemental Figure 4. Differences in interaction strength between resource-explicit models and the circle-intersection function, hexagonally-tiled. Two million pairwise interaction strengths between randomly placed individuals as measured by the circle-intersection function were subtracted from those measured in resource-explicit models to yield a distribution of the deviation from the circle-intersection function for each model. As node density increases, the standard deviation decreases. These distributions all have a distinctive peak just below 0 due to cases where pairs of individuals have a very small but non-zero interaction strength when using the circle-intersection function, but the small overlapping portion of their foraging areas does not include any resource nodes.

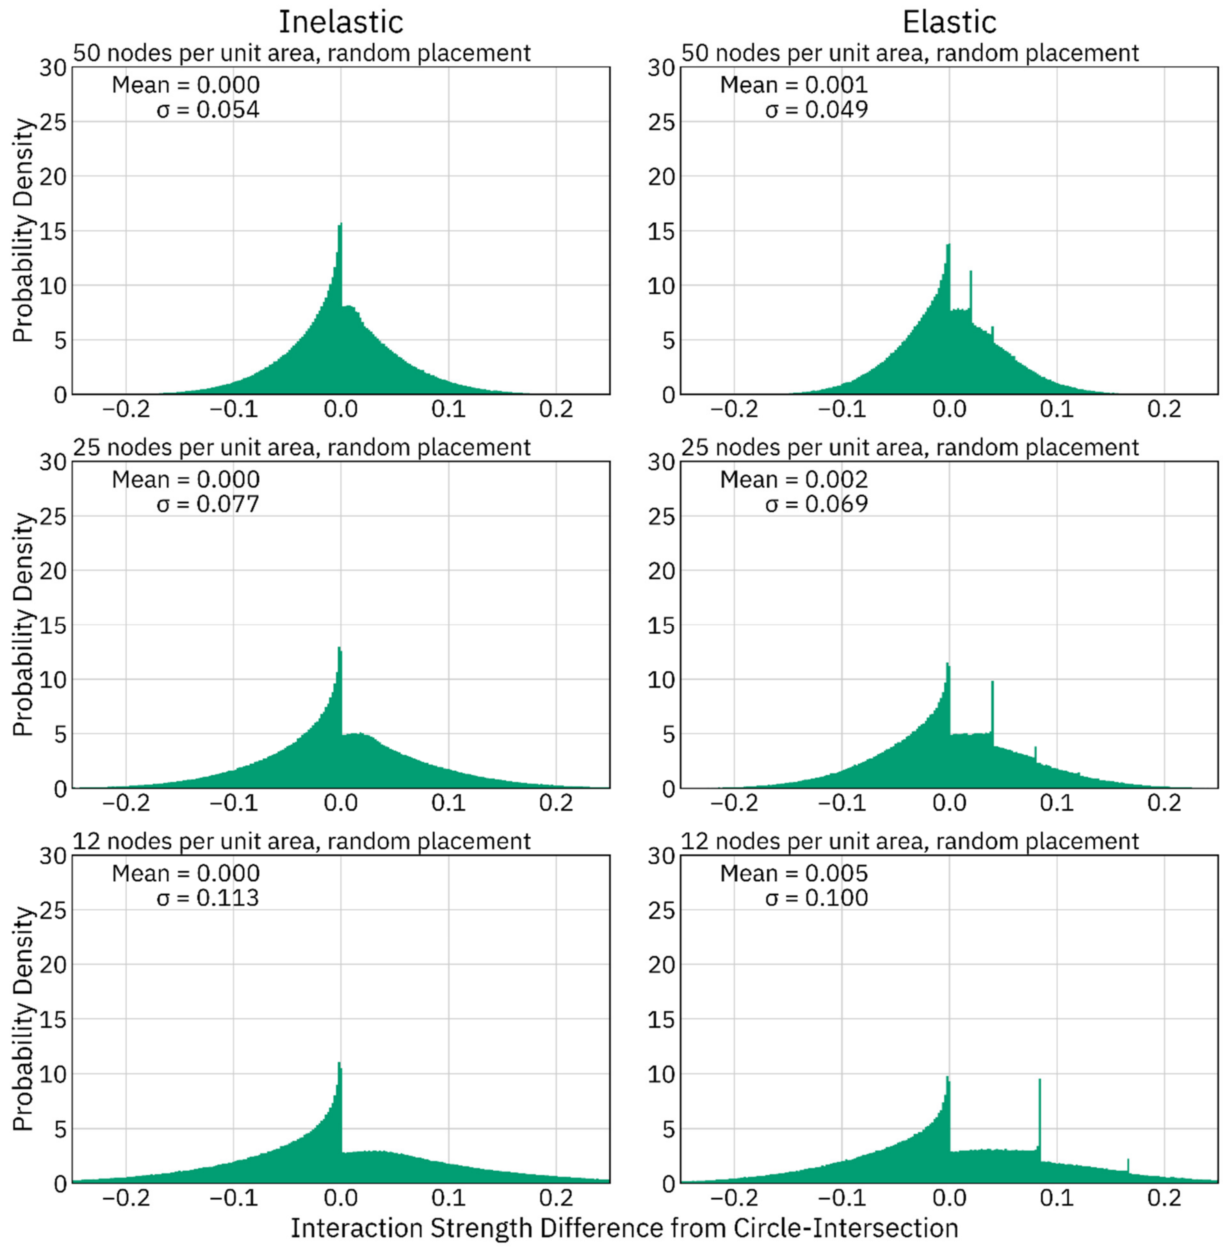

Supplemental Figure 5. Differences in interaction strength between resource-explicit models and the circle-intersection function, random node placement. Two million pairwise interaction strengths between randomly placed individuals as measured by the circle-intersection function were subtracted from those measured in resource-explicit models to yield a distribution of the deviation from the circle-intersection function for each model. As node density increases, the standard deviation decreases. These distributions all have a distinctive peak just below 0 due to cases where pairs of individuals have a very small but non-zero interaction strength when using the circle-intersection function, but the small overlapping portion of their foraging areas does not include any resource nodes. The distinctive peaks at positive values in the elastic model represent cases where the foraging circles do not intersect at all, but two individuals nonetheless share one or two nodes due to a lack of resource nodes closer to the individuals.

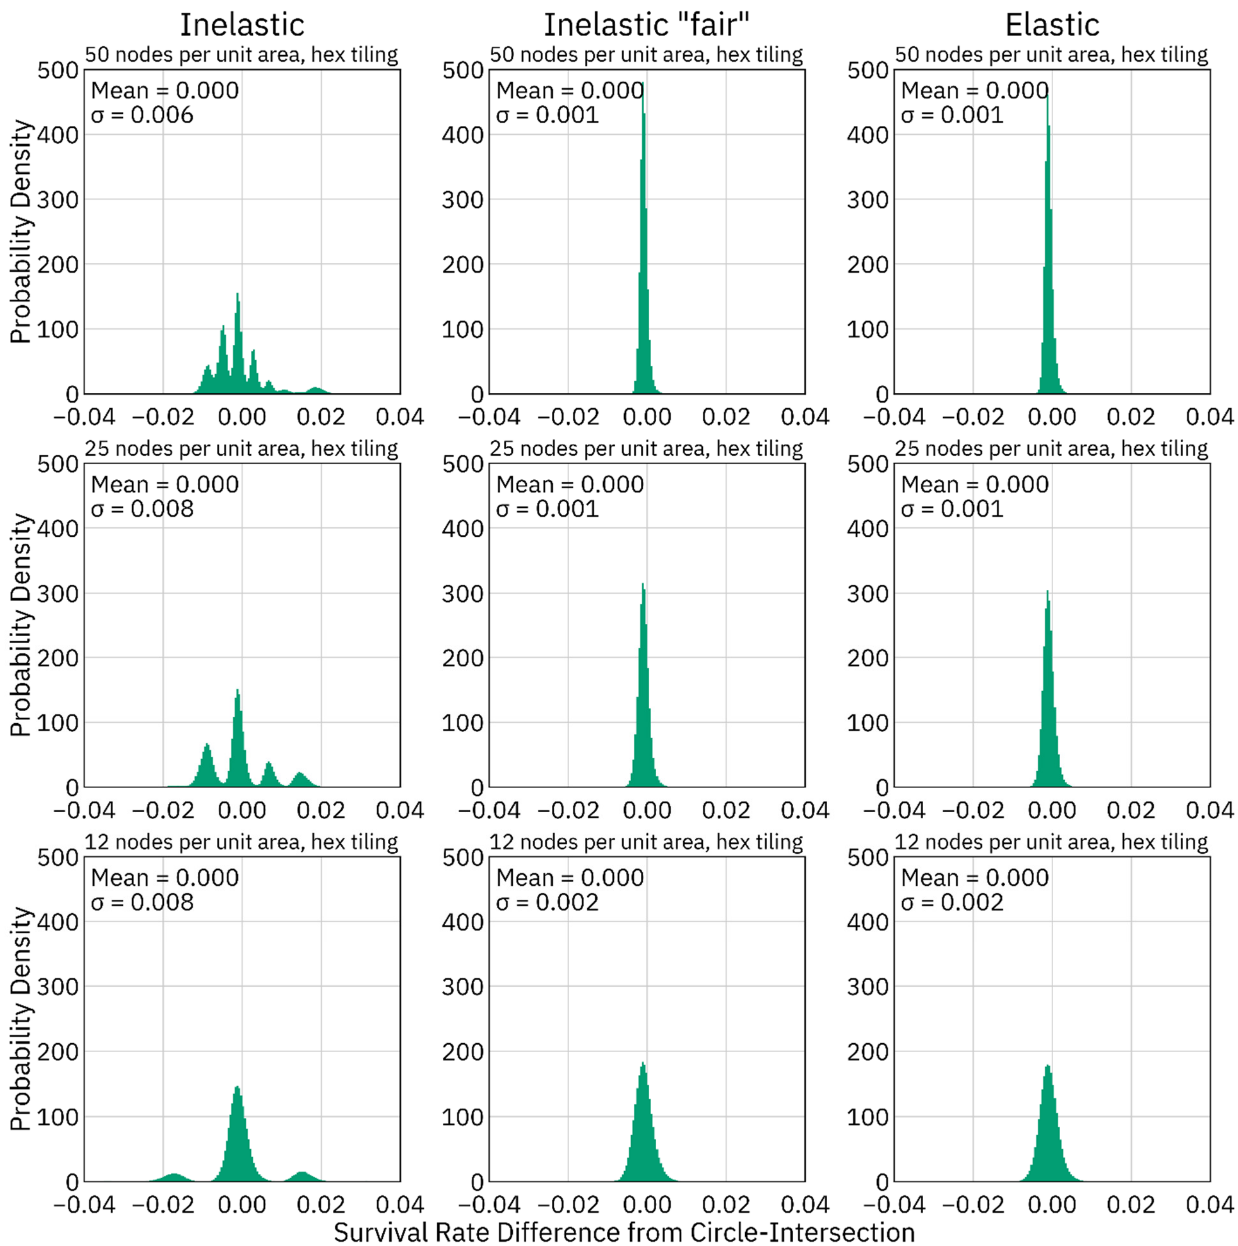

Supplemental Figure 6. Differences in survival rate between resource-explicit models and a direct-interaction model using the circle-intersection function, hexagonally-tiled. The survival rates of 100,000 individuals were measured using the inelastic method, the "fair" inelastic method, the elastic method, and the circle-intersection function. Strengths measured by the circle-intersection function were subtracted from measurements made using the other resource-explicit methods for each individual to yield distributions of differences.

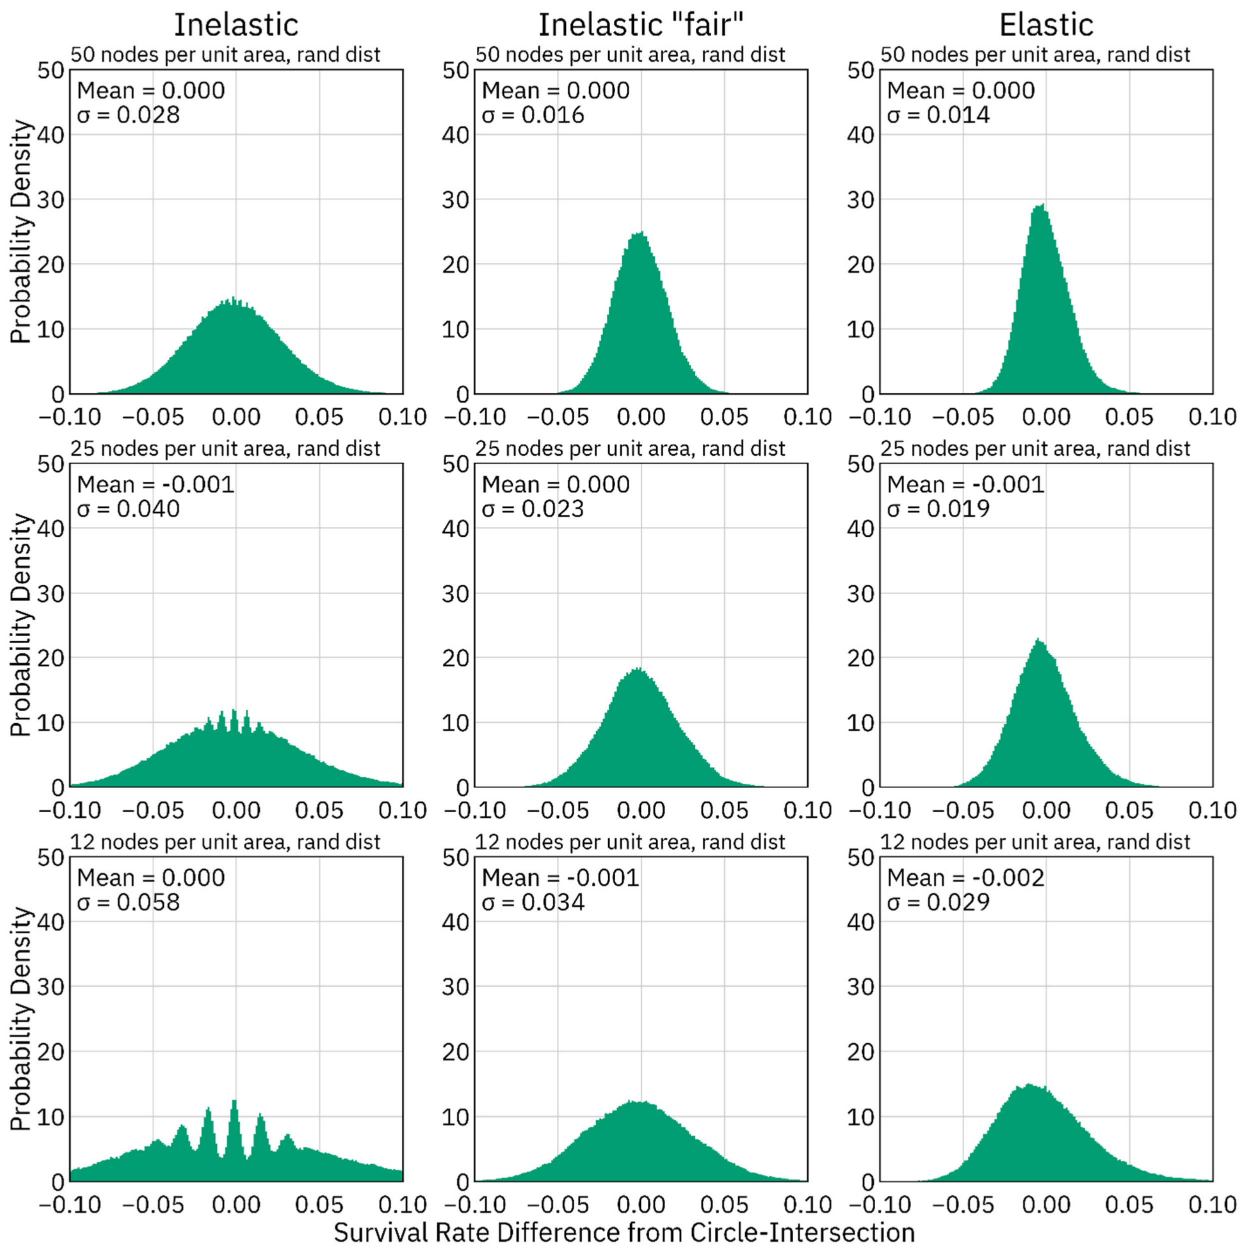

Supplemental Figure 7. Differences in survival rate between resource-explicit models and a direct-interaction model using the circle-intersection function, random node placement. The survival rates of 100,000 individuals were measured using the inelastic method, the "fair" inelastic method, the elastic method, and the circle-intersection function. Strengths measured by the circle-intersection function were subtracted from measurements made using the other resource-explicit methods for each individual to yield distributions of differences.

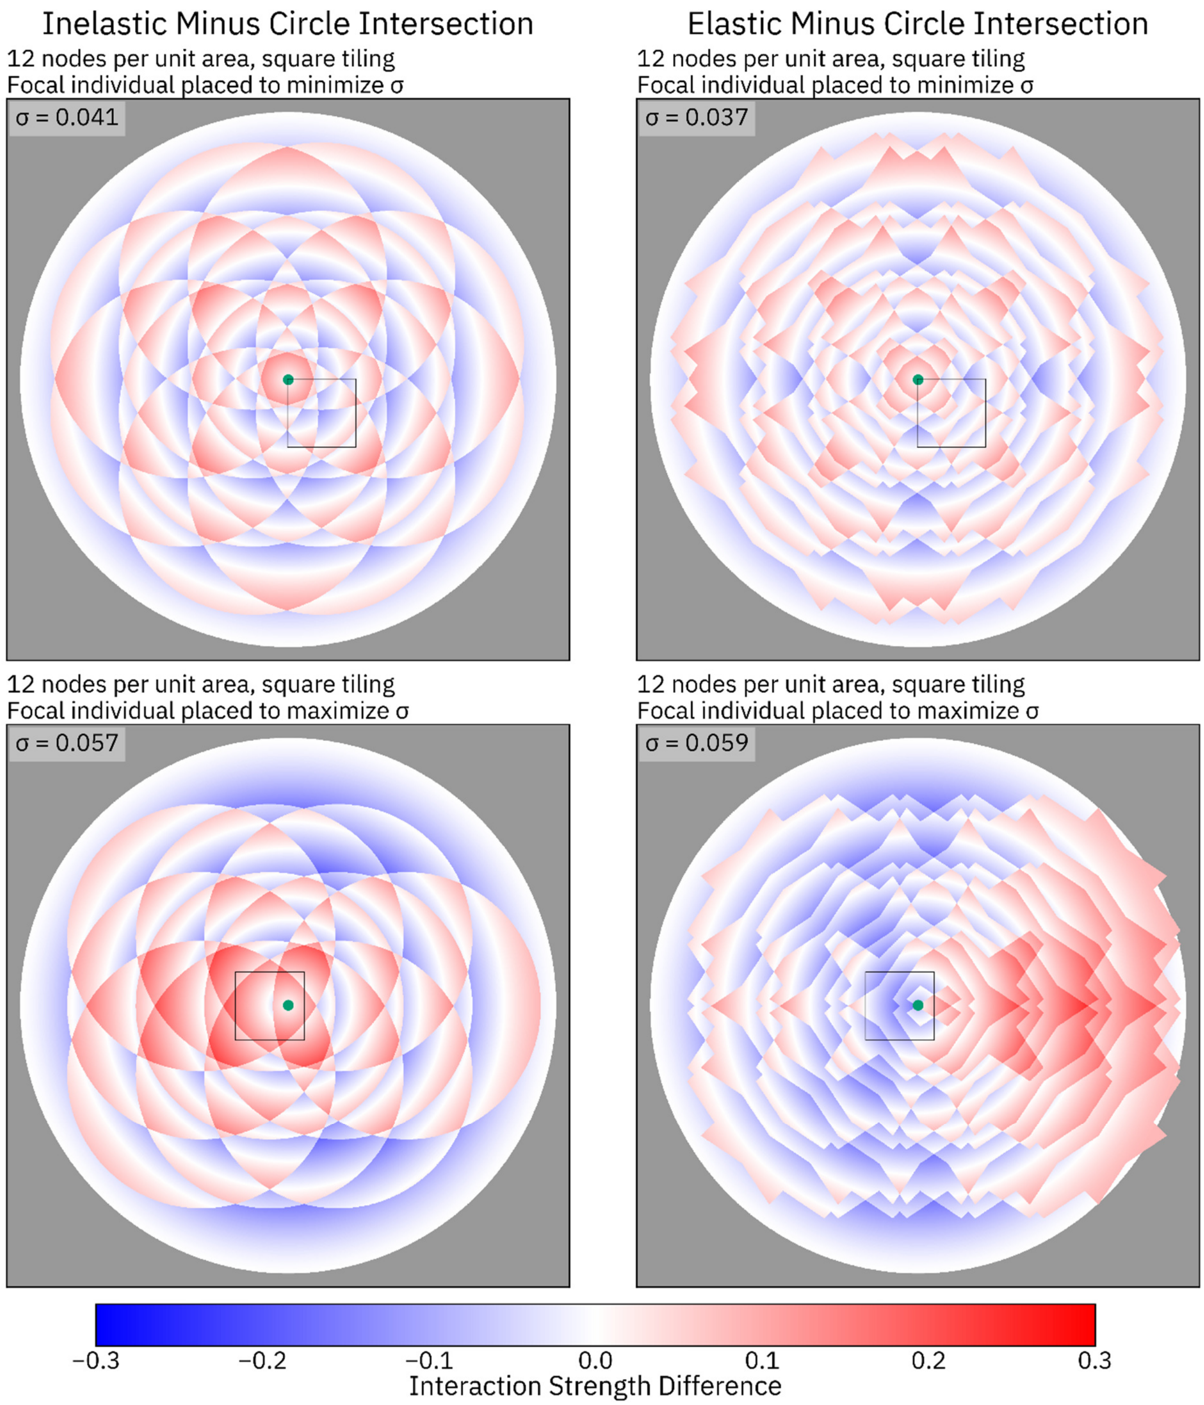

Supplemental Figure 8. Visualization of the resource-explicit interaction, square tiling, 12 nodes per unit area. A central individual (green dot) was placed to either minimize (top panels) or maximize (bottom panels) the standard deviation between the resource-explicit functions and the circle-intersection function. Interaction strength between the central individual and each other pixel of the image was measured using both a resource-explicit method and the circle-intersection function, and the difference is depicted. Blue shading indicates that the circle-intersection function is stronger, while red shading indicates that the resource-explicit interaction is stronger.

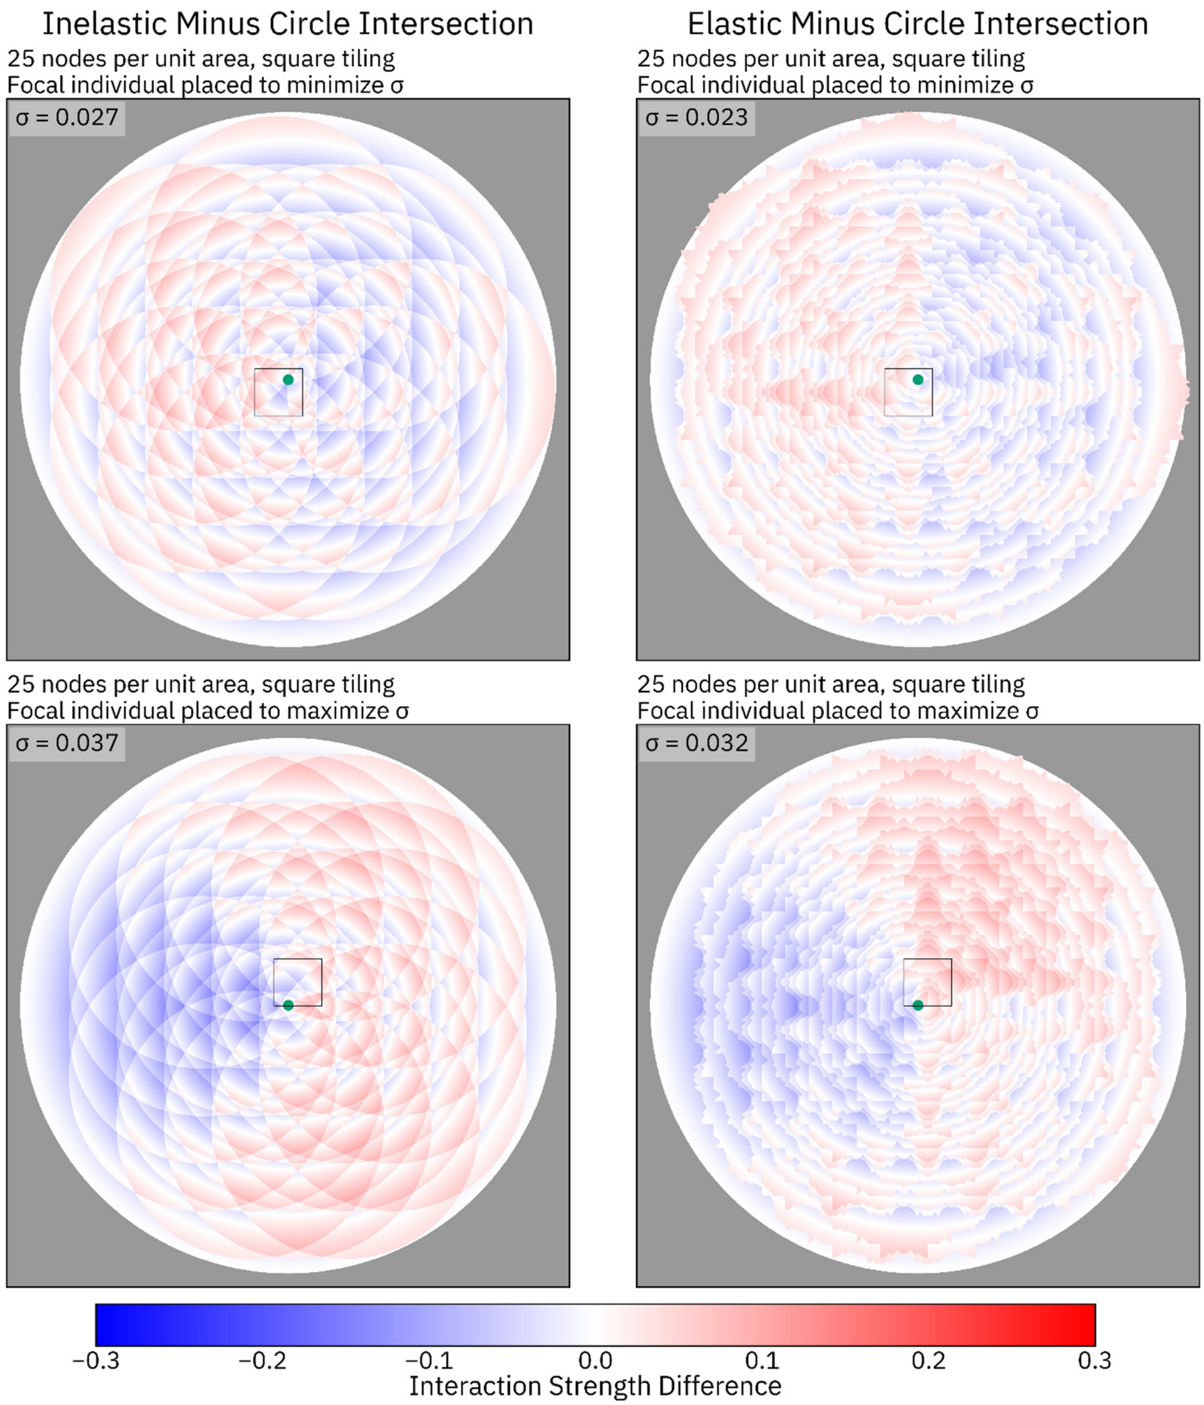

Supplemental Figure 9. Visualization of the resource-explicit interaction, square tiling, 25 nodes per unit area. A central individual (green dot) was placed to either minimize (top panels) or maximize (bottom panels) the standard deviation between the resource-explicit functions and the circle-intersection function. Interaction strength between the central individual and each other pixel of the image was measured using both a resource-explicit method and the circle-intersection function, and the difference is depicted. Blue shading indicates that the circle-intersection function is stronger, while red shading indicates that the resource-explicit interaction is stronger.

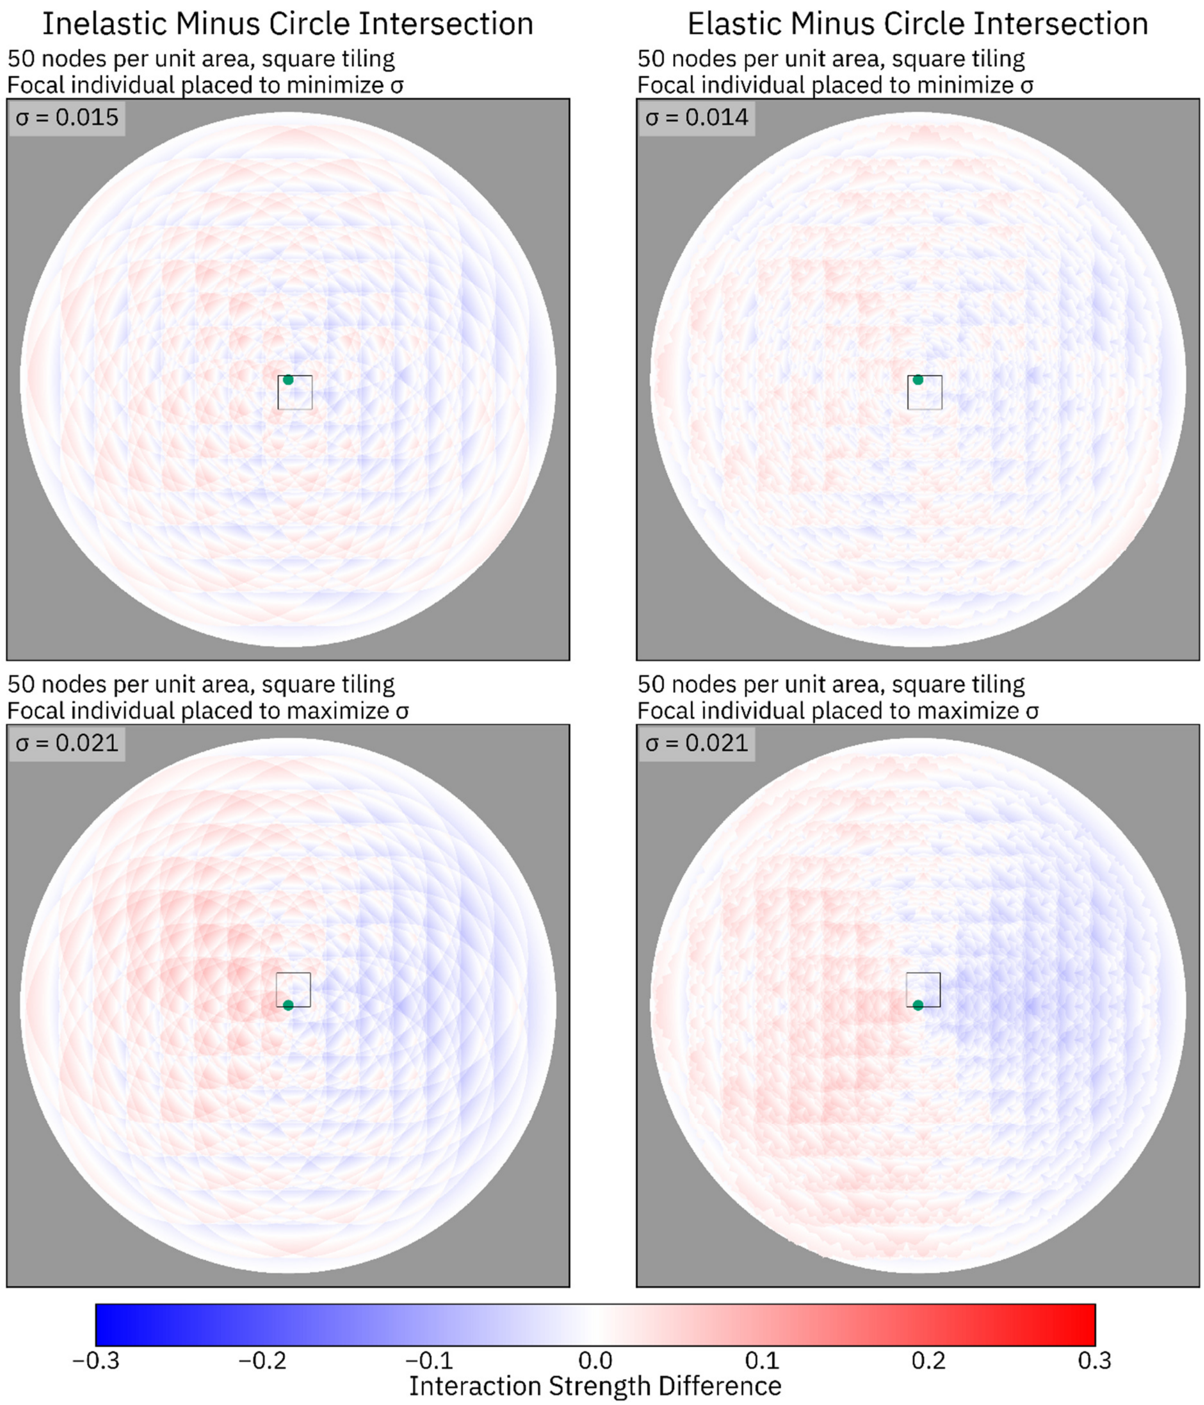

Supplemental Figure 10. Visualization of the resource-explicit interaction, square tiling, 50 nodes per unit area. A central individual (green dot) was placed to either minimize (top panels) or maximize (bottom panels) the standard deviation between the resource-explicit functions and the circle-intersection function. Interaction strength between the central individual and each other pixel of the image was measured using both a resource-explicit method and the circle-intersection function, and the difference is depicted. Blue shading indicates that the circle-intersection function is stronger, while red shading indicates that the resource-explicit interaction is stronger.

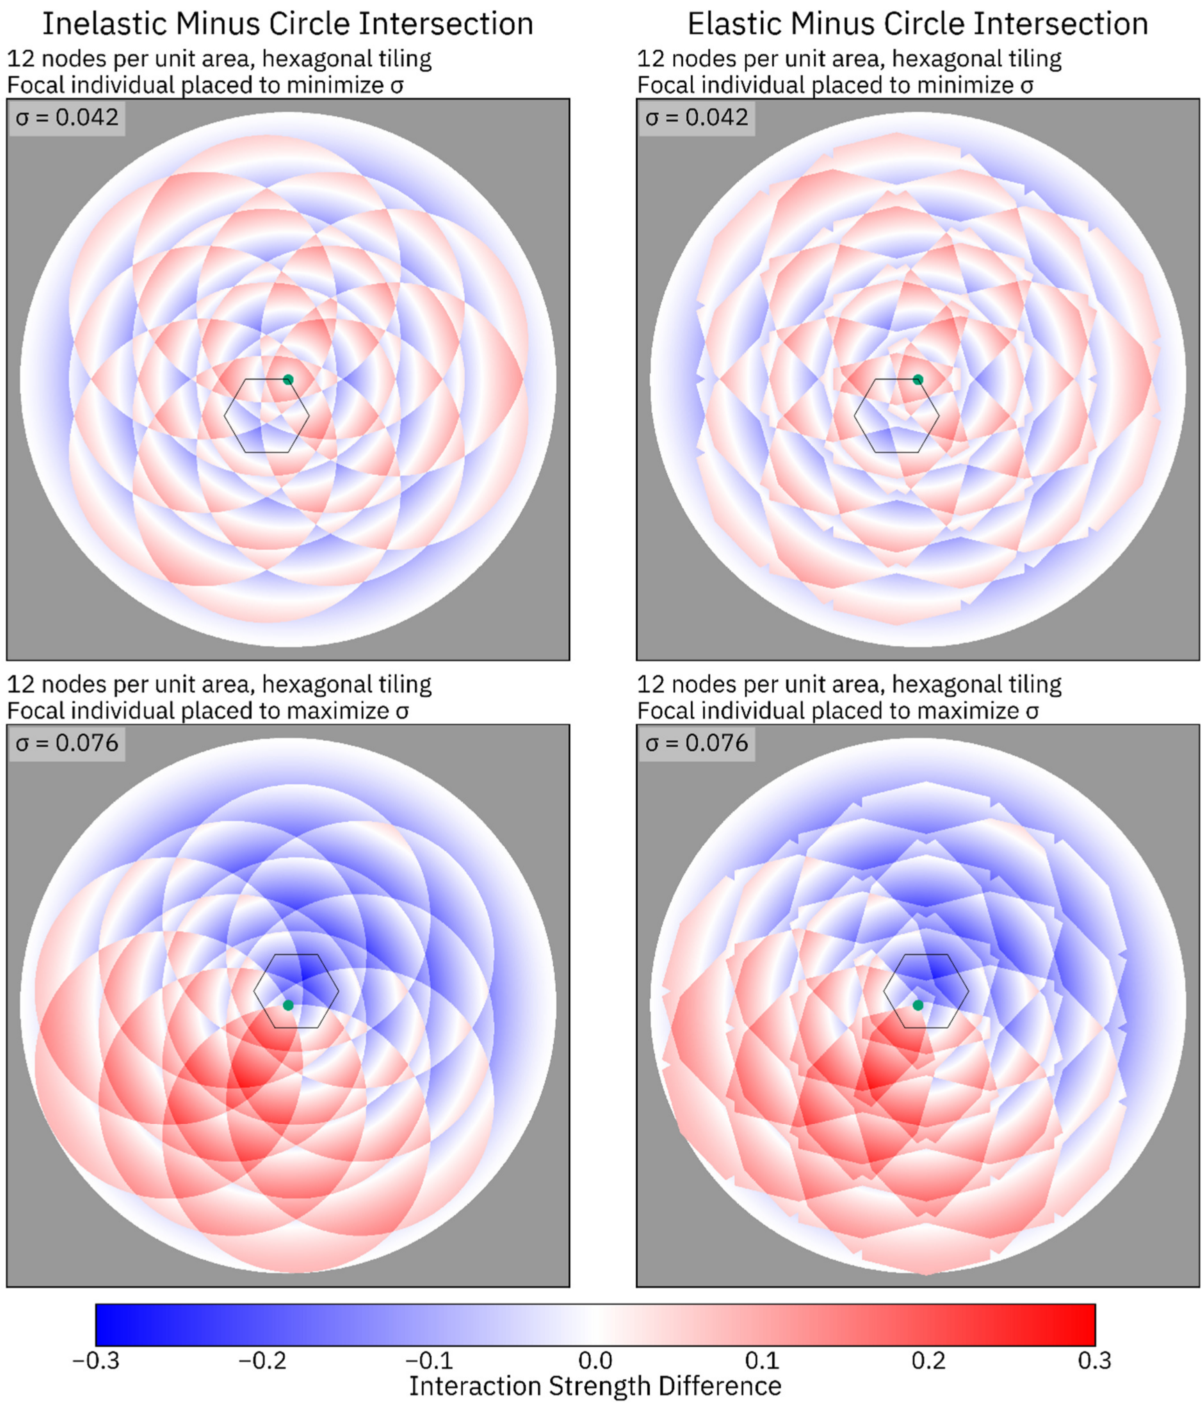

Supplemental Figure 11. Visualization of the resource-explicit interaction, hexagonal tiling, 12 nodes per unit area. A central individual (green dot) was placed to either minimize (top panels) or maximize (bottom panels) the standard deviation between the resource-explicit functions and the circle-intersection function. Interaction strength between the central individual and each other pixel of the image was measured using both a resource-explicit method and the circle-intersection function, and the difference is depicted. Blue shading indicates that the circle-intersection function is stronger, while red shading indicates that the resource-explicit interaction is stronger.

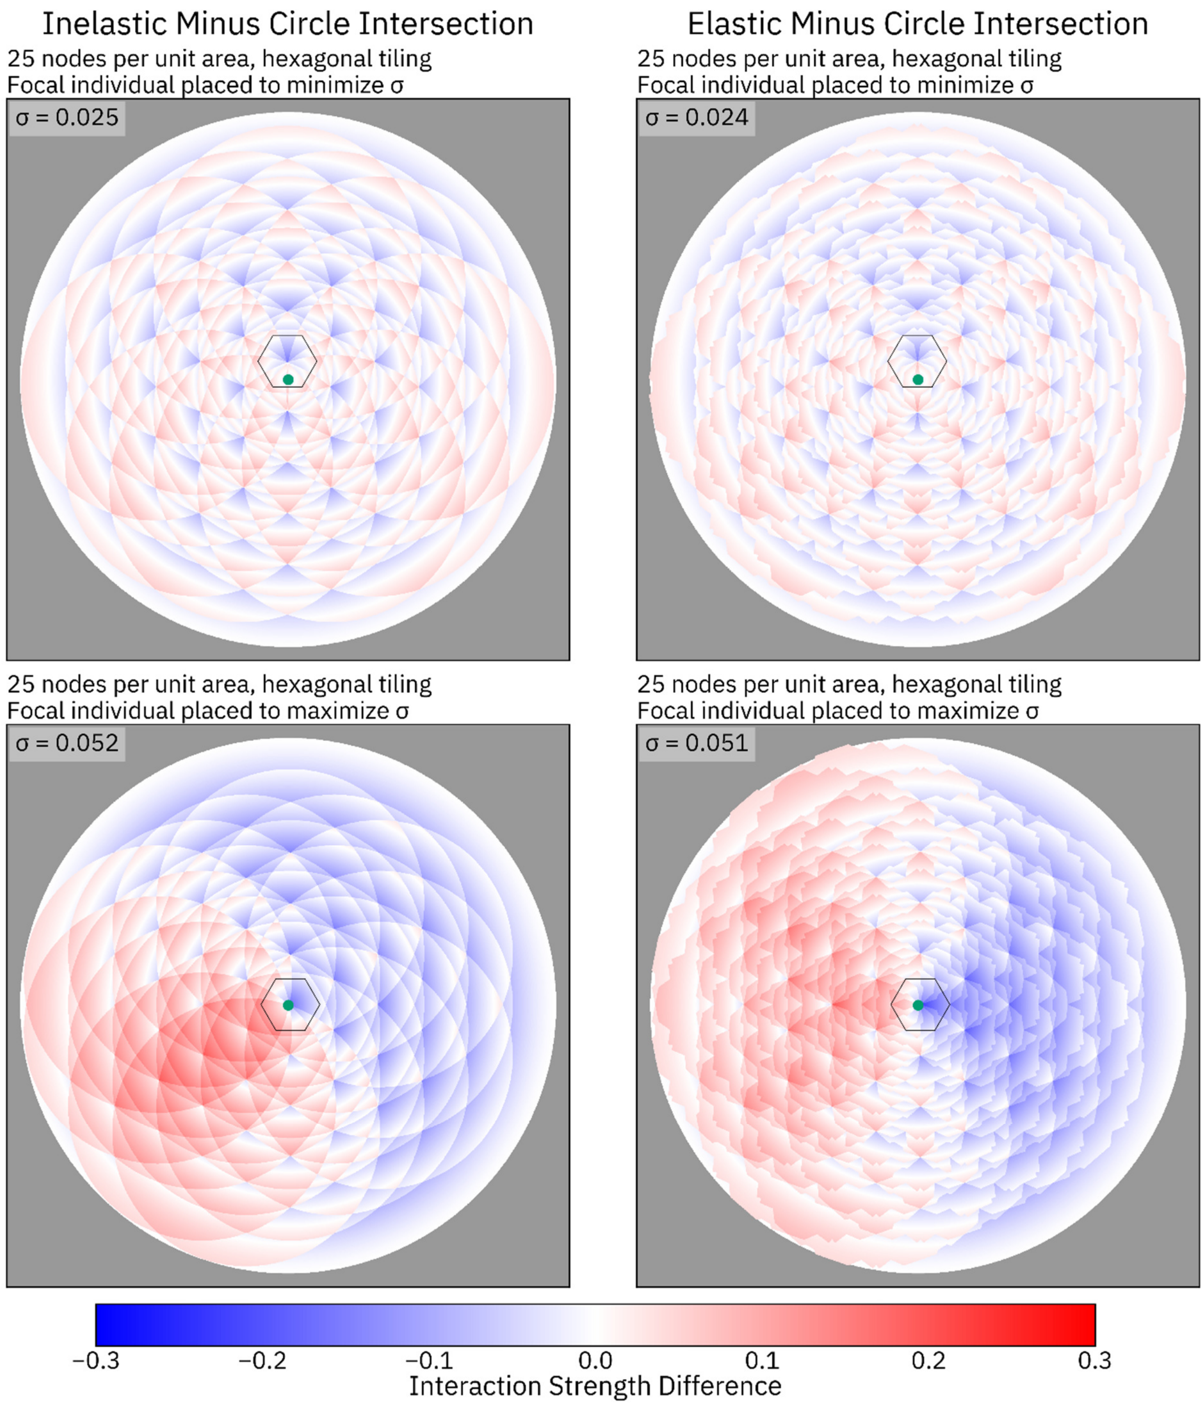

Supplemental Figure 12. Visualization of the resource-explicit interaction, hexagonal tiling, 25 nodes per unit area. A central individual (green dot) was placed to either minimize (top panels) or maximize (bottom panels) the standard deviation between the resource-explicit functions and the circle-intersection function. Interaction strength between the central individual and each other pixel of the image was measured using both a resource-explicit method and the circle-intersection function, and the difference is depicted. Blue shading indicates that the circle-intersection function is stronger, while red shading indicates that the resource-explicit interaction is stronger.

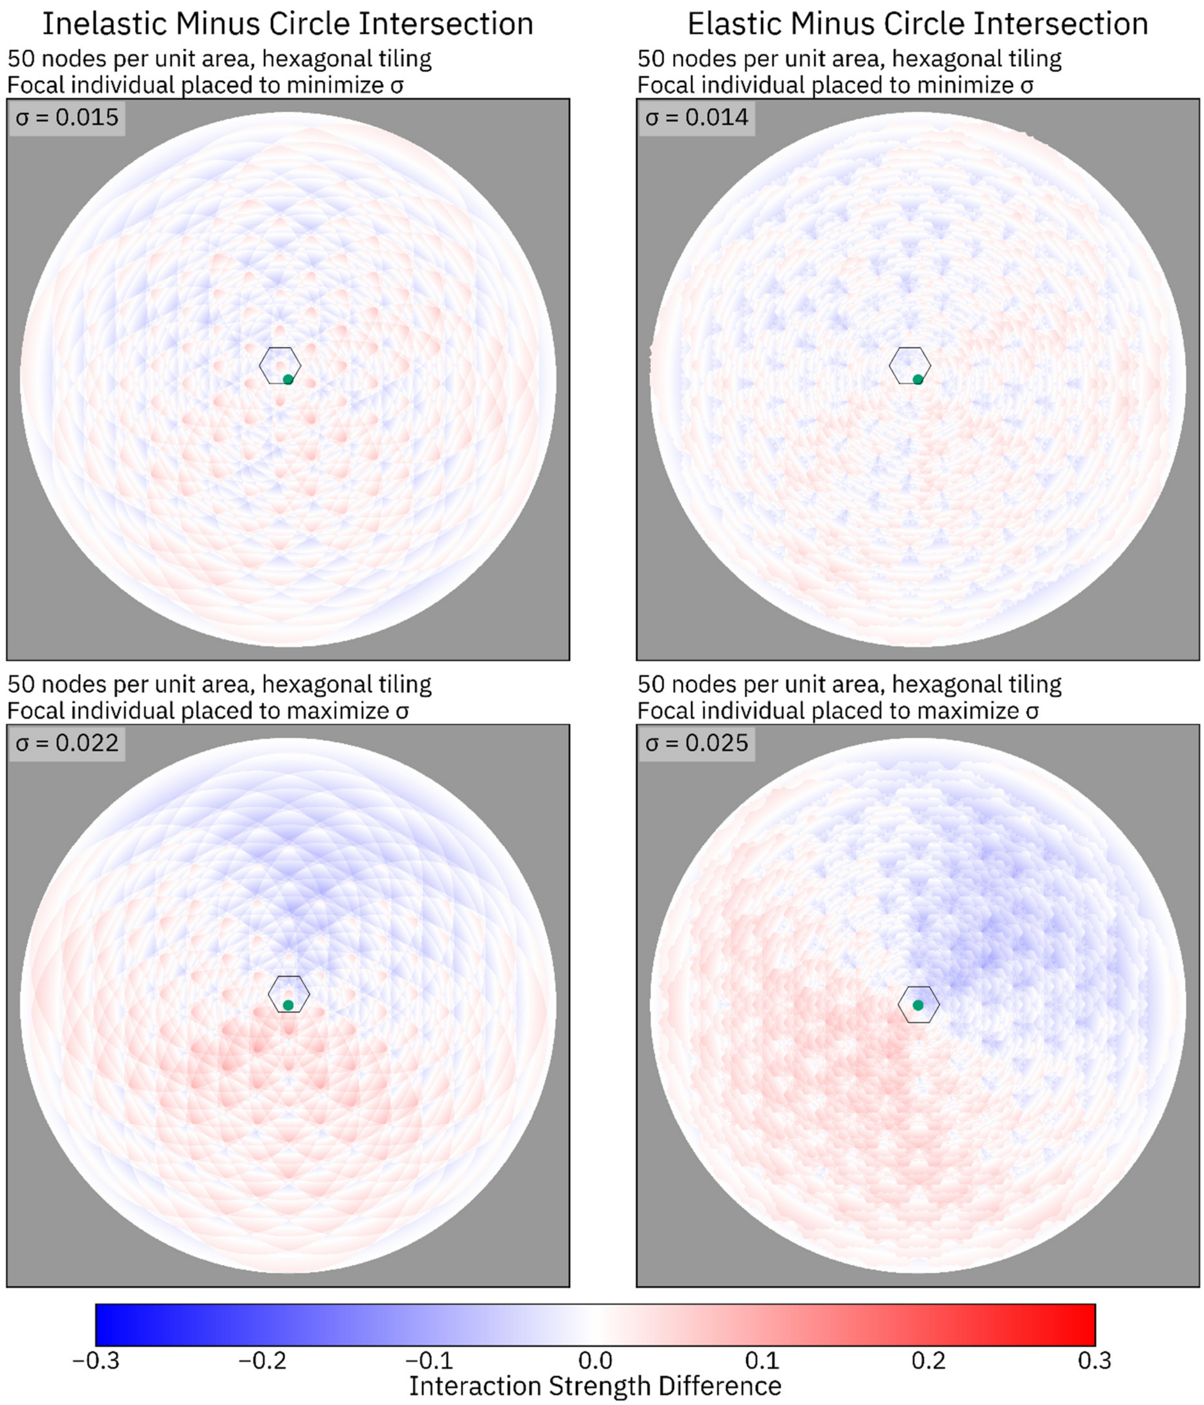

Supplemental Figure 13. Visualization of the resource-explicit interaction, hexagonal tiling, 50 nodes per unit area. A central individual (green dot) was placed to either minimize (top panels) or maximize (bottom panels) the standard deviation between the resource-explicit functions and the circle-intersection function. Interaction strength between the central individual and each other pixel of the image was measured using both a resource-explicit method and the circle-intersection function, and the difference is depicted by blue and red shading. Unlike in Supplemental Figs 8-12, the coordinate that maximized σ differs between the elastic and inelastic method in this case.

Unequal Resource Availability in the Inelastic Model. In the inelastic implementation of the method, individuals are not guaranteed to forage from the exact nominal foraging area – an individual might forage from more or fewer nodes. When the area is tiled with a uniform grid of nodes, the exact position of an individual relative to this tiling is what determines how many nodes fall within its foraging radius (Supplemental Figs 14-19).

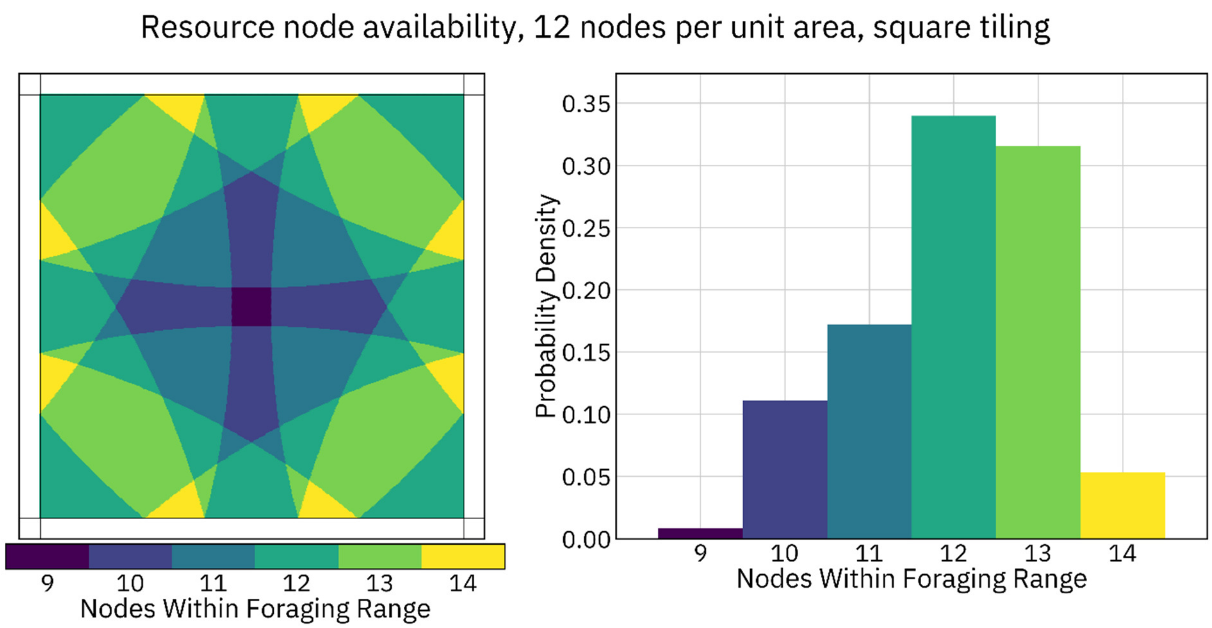

Supplemental Figure 14. Unequal access to resource nodes, square tiling, 12 nodes per unit area. The number of resource nodes within the foraging radius of an individual depends on their location within a square tile. Individuals located in darker regions of the square (left panel) forage from fewer nodes, while those in brighter regions forage from more. The relative area of each color is depicted in the right panel. The mean value of this distribution is 12.00, and the standard deviation is 1.11 (equivalent to 9.2 percent of the foraging area).

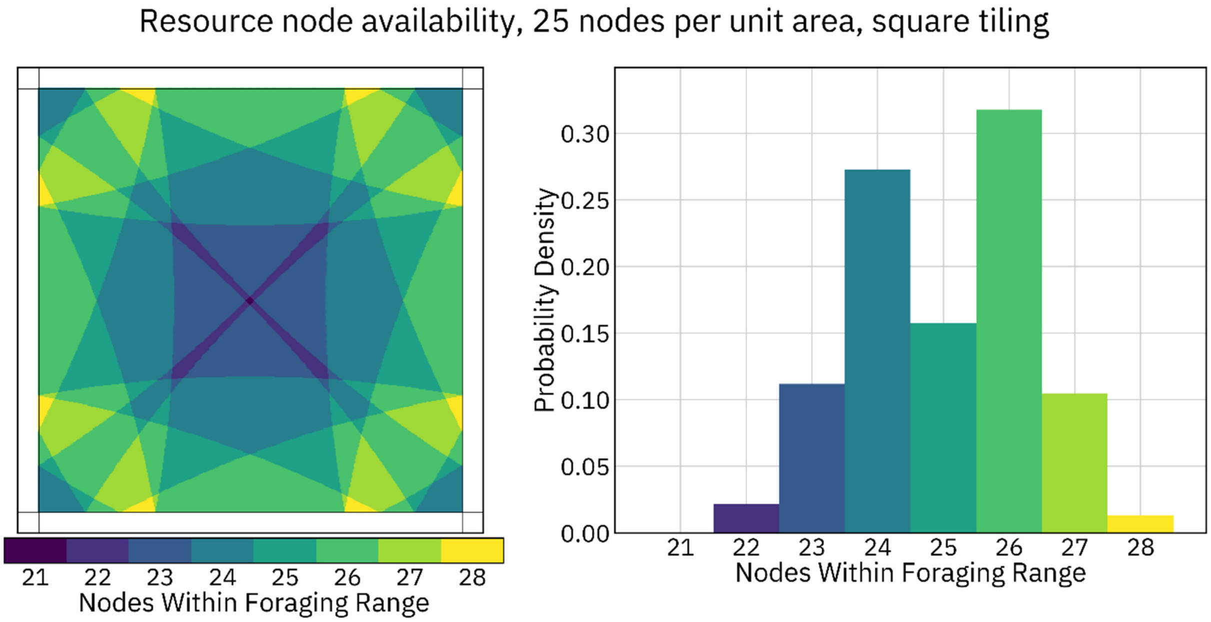

Supplemental Figure 15. Unequal access to resource nodes, square tiling, 25 nodes per unit area. The number of resource nodes within the foraging radius of an individual depends on their location within a square tile. Individuals located in darker regions of the square (left panel) forage from fewer nodes, while those in brighter regions forage from more. The relative area of each color is depicted in the right panel. The mean value of this distribution is 25.00, and the standard deviation is 1.33 (equivalent to 5.3 percent of the foraging area).

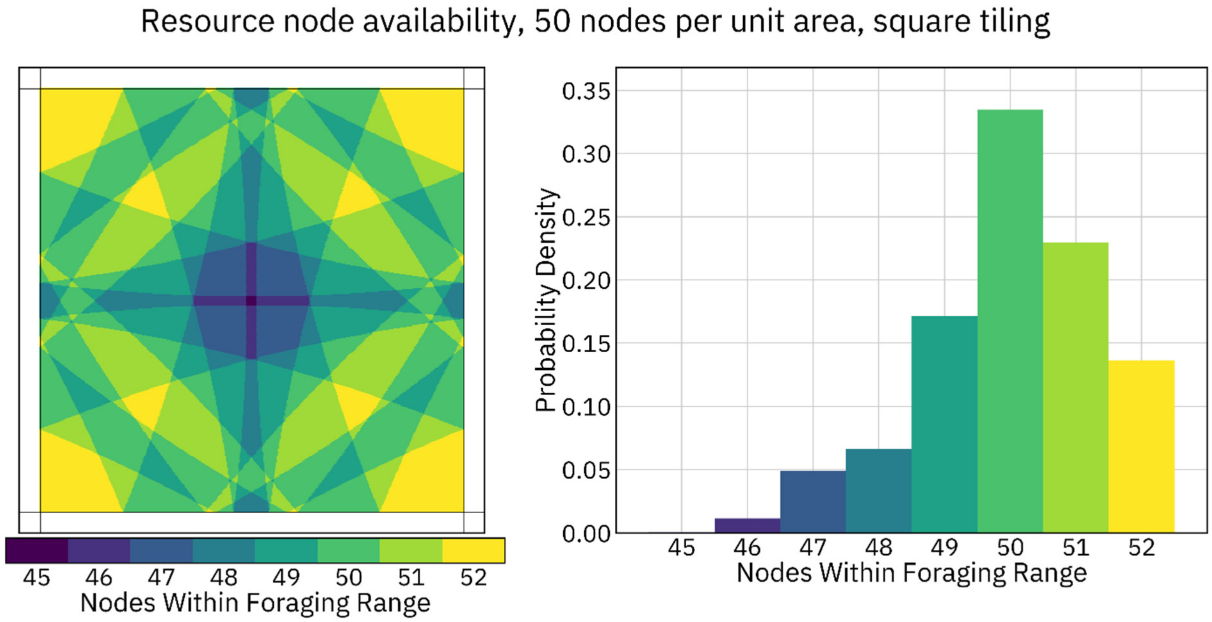

Supplemental Figure 16. Unequal access to resource nodes, square tiling, 50 nodes per unit area. The number of resource nodes within the foraging radius of an individual depends on their location within a square tile. Individuals located in darker regions of the square (left panel) forage from fewer nodes, while those in brighter regions forage from more. The relative area of each color is depicted in the right panel. The mean value of this distribution is 50.00, and the standard deviation is 1.36 (equivalent to 2.7 percent of the foraging area).

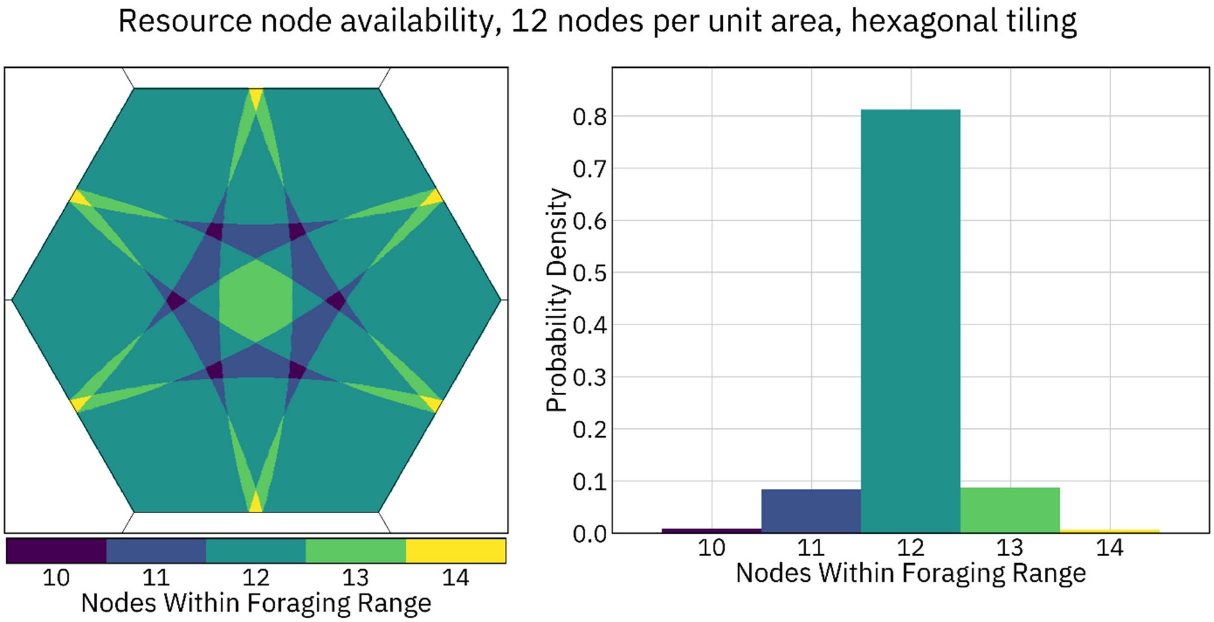

Supplemental Figure 17. Unequal access to resource nodes, hexagonal tiling, 12 nodes per unit area. The number of resource nodes within the foraging radius of an individual depends on their location within a hexagonal tile. Individuals located in darker regions of the hexagon (left panel) forage from fewer nodes, while those in brighter regions forager from more. The relative area of each color is depicted in the right panel. The mean value of this distribution is 12.00, and the standard deviation is 0.48 (equivalent to a 4.0 percent of the foraging area).

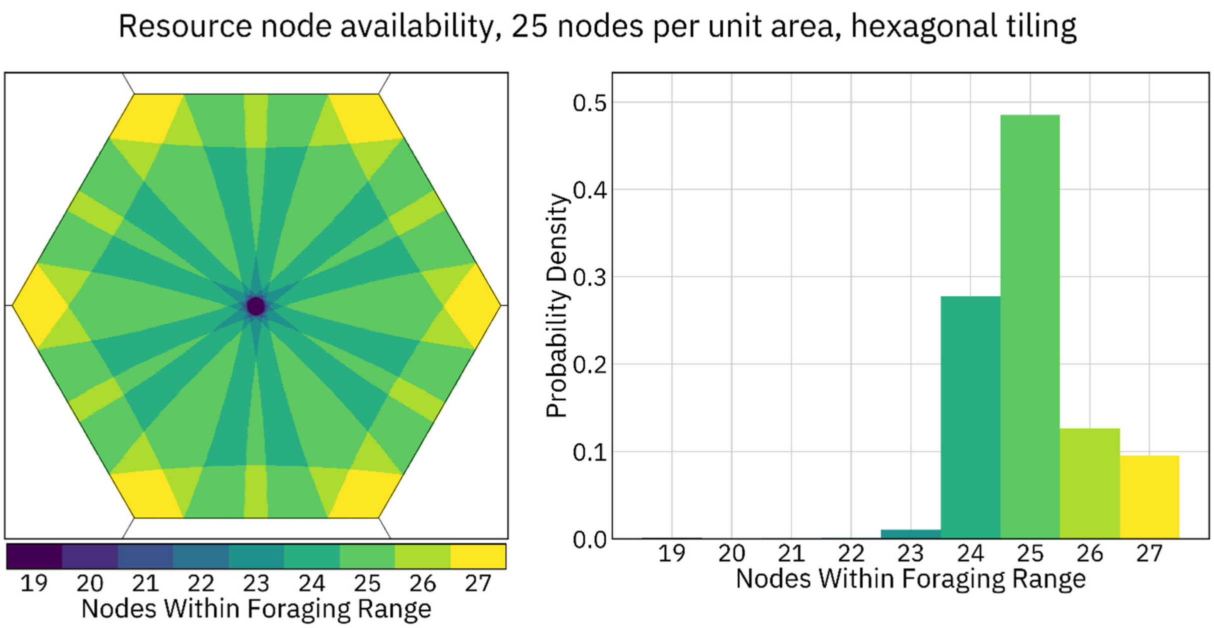

Supplemental Figure 18. Unequal access to resource nodes, hexagonal tiling, 25 nodes per unit area. The number of resource nodes within the foraging radius of an individual depends on their location within a hexagonal tile. Individuals located in darker regions of the hexagon (left panel) forage from fewer nodes, while those in brighter regions forager from more. The relative area of each color is depicted in the right panel. The mean value of this distribution is 25.00, and the standard deviation is 0.96 (equivalent to a 3.8 percent of the foraging area).

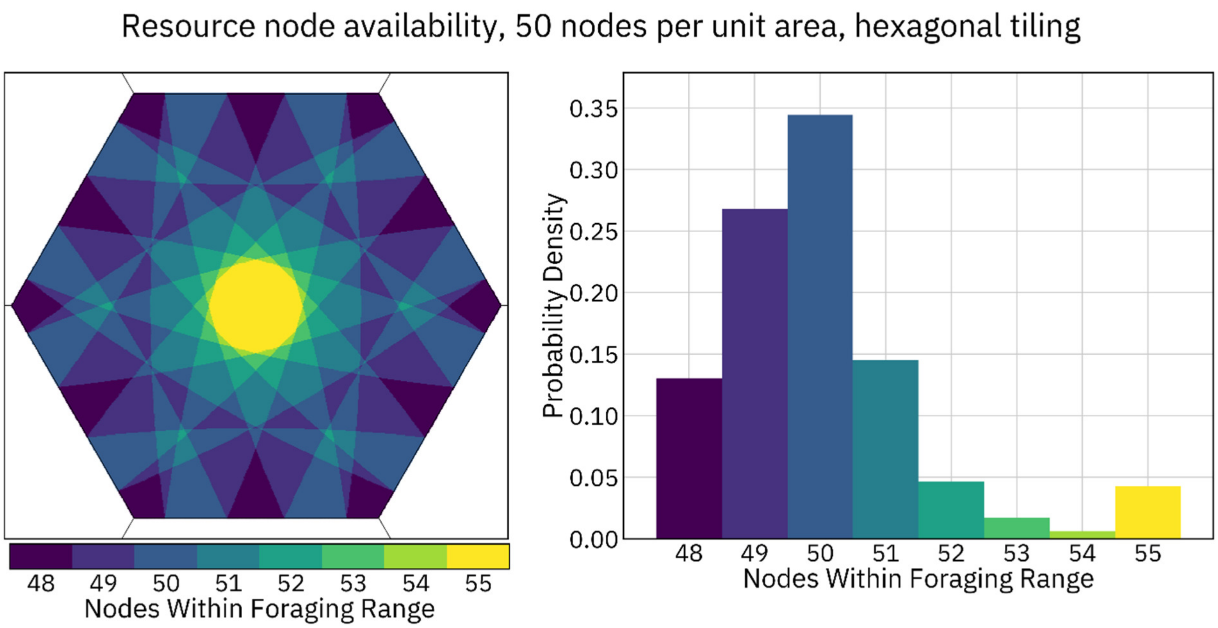

Supplemental Figure 19. Unequal access to resource nodes, hexagonal tiling, 50 nodes per unit area. The number of resource nodes within the foraging radius of an individual depends on their location within a hexagonal tile. Individuals located in darker regions of the hexagon (left panel) forage from fewer nodes, while those in brighter regions forager from more. The relative area of each color is depicted in the right panel. The mean value of this distribution is 50.00, and the standard deviation is 1.56 (equivalent to a 3.1 percent of the foraging area).

COMPUTATIONAL PERFORMANCE OF THE RESOURCE-EXPLICIT MODELS

The runtime differences between the square-tiled models and hexagon-tiled models were negligible, representing at most a small fraction of the runtime of the model (Supplemental Table 1). When resource nodes were randomly distributed during each tick of the model, runtimes were slower (Supplemental Table 2). Unsurprisingly, the models most slowed were those with the most resource nodes, which in some cases took almost twice as long to run as a uniformly tiled model with the same population size and node density.

Model	Node density	Low density (20 individuals per unit area)		High density (200 individuals per unit area)	
		100k individuals Runtime per tick (s)	1m individuals Runtime per tick (s)	100k individuals Runtime per tick (s)	1m individuals Runtime per tick (s)
Panmictic	NA	0.16	1.59	0.16	1.59
Fixed-strength	NA	3.03	38.98	17.01	233.43
Linear	NA	3.58	45.06	23.48	302.95
Gaussian	NA	5.76	67.88	44.81	523.23
Inelastic	12	0.83	9.42	1.18	15.71
	25	0.98	11.31	1.33	17.53
	50	1.33	15.05	1.57	21.27
Inelastic "fair"	12	1.00	11.77	1.35	18.16
	25	1.33	15.81	1.66	22.75
	50	2.04	23.50	2.26	31.15
Inelastic "resource-explicit reproduction"	12	0.98	10.95	0.82	10.59
	25	1.38	15.71	1.18	15.20
	50	2.14	24.40	1.91	24.10
Elastic	12	2.42	27.97	2.66	30.77
	25	3.51	40.95	3.63	41.42
	50	5.85	66.10	5.59	65.45

Supplemental Table 1. Model runtimes, hexagonal tiling. Forty measurements of elapsed runtime per tick were collected and averaged for each method at each density and population size and at each of three node-placement densities. For the resource-explicit models depicted in this table, a hexagonal tiling of nodes was used. Color bars show the runtime of each model, on a different scale within each column, relative to the slowest runtime within that column.

Model	Node density	Low density (20 individuals per unit area)		High density (200 individuals per unit area)	
		100k individuals Runtime per tick (s)	1m individuals Runtime per tick (s)	100k individuals Runtime per tick (s)	1m individuals Runtime per tick (s)
Panmictic	NA	0.16	1.59	0.16	1.59
Fixed-strength	NA	3.03	38.98	17.01	233.43
Linear	NA	3.58	45.06	23.48	302.95
Gaussian	NA	5.76	67.88	44.81	523.23
Inelastic	12	0.95	11.94	1.27	16.77
	25	1.29	17.10	1.50	20.32
	50	1.97	27.03	1.92	26.57
Inelastic "fair"	12	1.23	15.44	1.51	19.65
	25	1.91	24.05	2.00	26.11
	50	3.23	41.09	2.98	38.72
Inelastic "resource-explicit reproduction"	12	1.19	14.67	0.92	12.40
	25	1.88	24.37	1.38	19.27
	50	3.15	44.24	2.41	32.62
Elastic	12	2.47	27.73	2.70	30.73
	25	3.52	41.58	3.63	42.08
	50	5.92	67.22	5.56	65.61

Supplemental Table 2. Model runtimes, random node placement. Forty measurements of elapsed runtime per tick were collected and averaged for each method at each density and population size and at each of three node-placement densities. For the resource-explicit models depicted in this table, nodes were randomly distributed across the landscape during each tick of the model. Color bars show the runtime of each model, on a different scale within each column, relative to the slowest runtime within that column.

ADDITIONAL EXTENSIONS

I. An Implementation of the Elastic Method Optimized for Infrequent Dispersal

An improvement in runtime can be achieved in the elastic model if each individual only disperses infrequently or only a single time (their initial dispersal from their maternal parent), by finding the nearest resource nodes to an individual only after it disperses, and then caching and reusing this list of nodes during each subsequent tick until the individual disperses again or dies. Individuals disperse only once in all of the models presented in this manuscript, but this optimization was not made in the default elastic model in order to provide an accurate reflection of the runtime that could be expected in a model where individuals move during each tick. This optimization is not relevant in a model with non-overlapping generations or in a model with nodes that are re-randomized each time step, since the cached node lists would never be reused.

The degree of runtime improvement that this optimization could yield is highly dependent on the demography of the model. In models with a low per-tick mortality and in which individuals only disperse once, this method could be faster than all of the other resource-explicit methods presented in this manuscript. In the model assessed in this study, the mortality rate was quite high (about 80 percent per tick) due to the large number of offspring produced every tick of the model, and the cached node lists only save time for individuals that survive to the next tick to use that cache. Thus, when modified to include this optimization, the elastic model used in this study only increased in speed by a small amount, and was still much slower than the inelastic model. The full code for this modification is provided on GitHub.

II. A Resource-Explicit Model with a Semi-Fixed Population Size

Many analytical models consider populations to consist of a fixed number of individuals. When extending pre-existing analytical models into an individual-based spatial context, it can be desirable to otherwise match the analytical model as closely as possible by maintaining a fixed population size in the spatial model (recognizing that this is biologically unrealistic). A modification of the resource-explicit modeling technique allows for a population size that is fixed except in the event of extreme disruptions to the population.

In implementations described in the Methods section, each node is parameterized with some amount of resources that depends on the density of the modeled species and the density of the resource nodes, and each node distributes its resources evenly to all individuals that forage from it. If, for example, there are 50 individuals foraging from a node with 10 resources available, each individual will receive 0.2 resources, contributing a 20 percent probability of survival to each of those individuals.

In this variant of the model with a semi-fixed population size, resource nodes are instead parameterized by an integer number of “tickets”. Instead of receiving a floating-point amount of resources, individuals instead have a chance to receive a ticket. After tickets are distributed, only individuals who have received a ticket survive. This results in the landscape maintaining a population of exactly the specified carrying capacity, except in the event of a disturbance in the model that causes reproduction to be insufficient to reach that carrying capacity (such as a simulated intervention against a pest population). Note that the population will also not stay exactly at a defined capacity in a model with a marginal habitat quality or a very low birthrate, in which case reproduction may sometimes produce fewer individuals than mortality removes even when no external factors are present.

For each of the four variants of the resource-explicit model presented in this manuscript (elastic, inelastic, inelastic “fair”, and inelastic with “resource-explicit reproduction”), the GitHub file repository for this project also contains an equivalent model that has been modified as described above to maintain a semi-fixed population size. This modification results in a moderate performance reduction in the elastic models, but does not appear to noticeably affect the performance of the inelastic models.

For modelers seeking to design spatial models that match analytical models as closely as possible, we anticipate that the “fair” variant of the inelastic model may be the best option. This model is faster than the elastic model, yet it avoids the small-scale spatial artifacts present in the default inelastic model which may be undesirable in this context (see Fig. 4 and Supplemental Figs 14-19).

III. A Spatial Model of the South Island of New Zealand

The resource-explicit method has the potential to scale up to modeling large populations in large heterogeneous habitats while maintaining relatively performant runtimes. As an initial exploration into this possibility, we implemented a model on a landscape map of the South Island of New Zealand. Endemic diversity in New Zealand is threatened by the presence of numerous invasive mammal species, and detailed spatial modeling is the first step in investigating potential population control strategies, such as gene drive, in order to maintain and restore biodiversity (Champer, Oakes, et al., 2021).

To produce a heterogeneous landscape map of the island, we constructed a habitat suitability map in which we defined habitat quality as a function of elevation, with optimum habitat at about 300 meters above sea level, with quality decreasing at higher and lower elevations. Each resource node was parameterized with resources according to the local elevation near that node. The total area of the South Island is 150,416 km^2, but mountainous regions were considered to have no accessible resources (Statistics New Zealand, 2010). The amount of the landscape defined as usable habitat by the generic focal species in the model was 100,636 km^2. The nominal foraging area of the focal species was defined as 0.25 km^2, and the landscape was populated with nodes using a square tiling at a density of 12 nodes per 0.25 km^2. The total number of resource nodes is 4.8 million. Ticks in the model represent monthly intervals. We included seasonality in the model as a per-tick multiplier to the resource value of each node that follows a sinusoidal function with a maximum of 1 in the summer and 0.5 in the winter. The carrying capacity of the focal species across the entire area is 10 million individuals in the summer, with a maximum density of about 50 individuals per 0.25 km^2 and an average density of about 25 individuals per 0.25 km^2 (and half these densities in the winter). Node coordinates and resource values were determined in a pre-processing step that generated a CSV file that was reused by all of the replicates of the simulation.

The inelastic implementation of the resource-explicit model was chosen for this simulation, in the interest of maximizing the speed of the model. Other than loading node positions from the external CSV file at the outset of the model, the only change to the model was to ensure that individuals were positioned on the landmass at the start of the simulation and were prevented from dispersing into the ocean.

No in-depth analyses were conducted of this model, nor were any analogous models constructed for comparison purposes. The average runtime of each tick of this model was under 1.5 minutes even in the summer, when the population was at its maximum size. This is almost certainly sufficiently fast to be used in a study, unless a very large number of ticks needs to be simulated. Though the dynamics within the model were not thoroughly analyzed, a visual assessment indicates that the heterogeneity of the landscape is satisfactorily reflected in the distribution of the population (Supplemental Fig. 20).

We also anecdotally observed a number of features of this model that were not explicitly programmed, but which are present as emergent properties of the resource availability of the landscape. These features include landscape fragmentation, with some populations on the landscape (not just the smaller islands) appearing to be completely separated from the main population. We also observed cases of partial fragmentation, wherein some inland populations were separated from a larger coastal population only during the winter, but were once again connected during the summer. We also observed source-sink dynamics in which a population in a marginal habitat area could continue to exist thanks to immigration from higher-quality habitat. This was confirmed by artificially removing individuals from some of the coastal areas of the model, after which some populations further inland collapsed due to marginal habitat quality combined with a lack of immigration from the adjacent (artificially vacated) coastal areas.

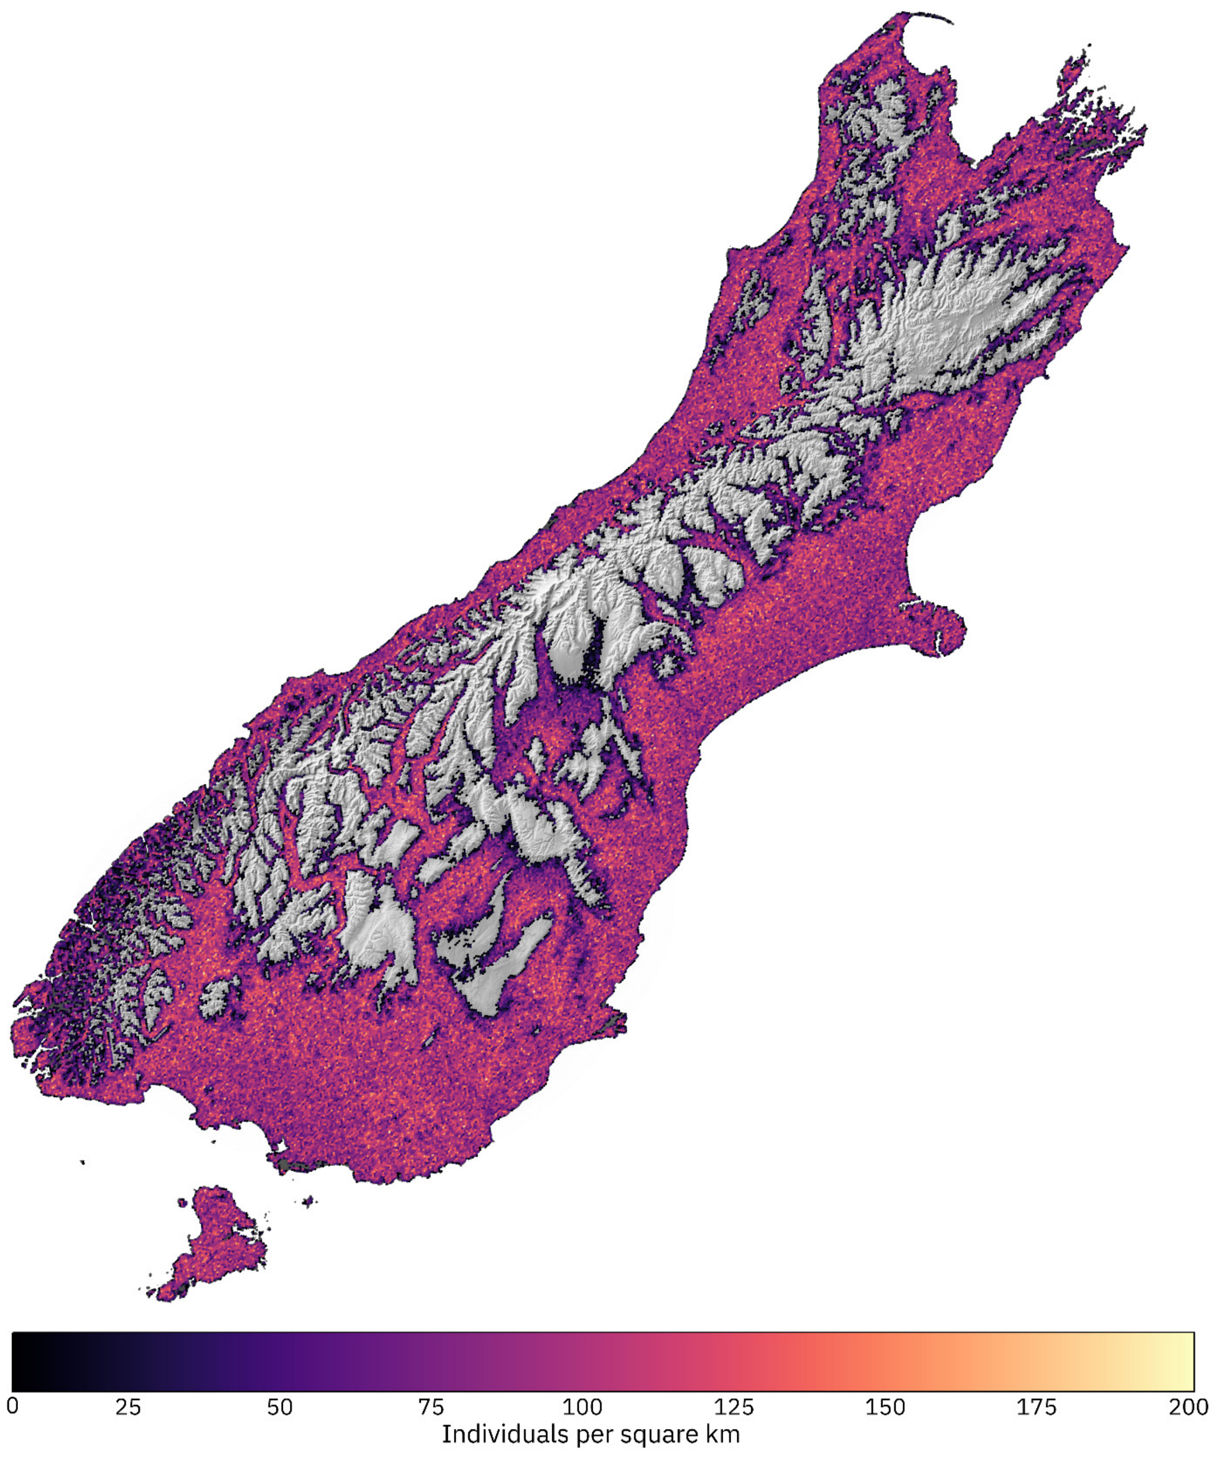

Supplemental Figure 20. A heterogeneous landscape model of New Zealand. The modeled area was populated with a square grid of 4.8 million resource nodes and 10 million individuals. Habitat quality was defined as loosely inversely proportional to altitude, with the optimum habitat at about 300 meters above sea level. The color shade gradient denotes the density of individuals. Topographical map image (visible as grayscale shading in areas with no resource nodes) courtesy NASA JPL.