Supplementary Table 1. **The ten focal annual plant species**

Species	Family	Life form	Rooting depth (cm)	Julian date when 50% of seeds have dehisced (2007)	Seed bank annual survival
Chaenactis glabriuscula	Asteraceae	Forb	81	134	68%
Chorizanthe palmeri	Polygonaceae	Forb	124	223	36%
Lasthenia californica	Asteraceae	Forb	56	104	31%
Lotus wrangelianus	Fabaceae	Legume	103	140	51%
Micropus californicus	Asteraceae	Forb	74	120	16%
Navarretia atractyloides	Polemoniaceae	Forb	110	235	16%
Plantago erecta	Plantaginaceae	Forb	104	133	59%
Salvia columbariae	Lamiaceae	Forb	169	147	28%
Trifolium willdenovii	Fabaceae	Legume	126	132	53%
Vulpia microstachys	Poaceae	Grass	50	131	21%

Supplementary Table 2. Model fit relative to seven other candidate models

Model	AIC (Akaike Information Criterion)
$F_i = \dfrac{\lambda_i}{1+\alpha_{ii} g_i N_i + \alpha_{ij} g_j N_j}$	**5203**
$F_i = \dfrac{\lambda_i}{1+\left(\alpha_{ii} g_i N_i + \alpha_{ij} g_j N_j\right)^{b_i}}$	5246
$F_i = \lambda_i e^{-\alpha_{ii} g_i N_i - \alpha_{ij} g_j N_j}$	5550
$F_i = \lambda_i e^{-\alpha_{ii} \ln\left(g_i N_i +1\right) - \alpha_{ij} \ln\left(g_j N_j +1\right)}$	5261
$F_i = \dfrac{\lambda_i}{1+\left(g_i N_i\right)^{\alpha_{ii}} + \left(g_j N_j\right)^{\alpha_{ij}}}$	5394
$F_i = \dfrac{\lambda_i}{\left(1+\alpha_{ii} g_i N_i + \alpha_{ij} g_j N_j\right)^{b_i}}$	5398
$F_i = 1+\lambda_i\left(1-\alpha_{ii} g_i N_i - \alpha_{ij} g_j N_j\right)$	6008
$F_i = \lambda_i$	6484

Maximum likelihood tests compared candidate models for how the number of viable seeds produced per germinant, F_i, declines with germinant density. The model in our study (see “Theoretical approach” in the Methods) had the lowest (best) AIC. Parameters are defined in the Methods, and the candidate models were all those presented in reference 1. See Supplementary methods for analysis details.

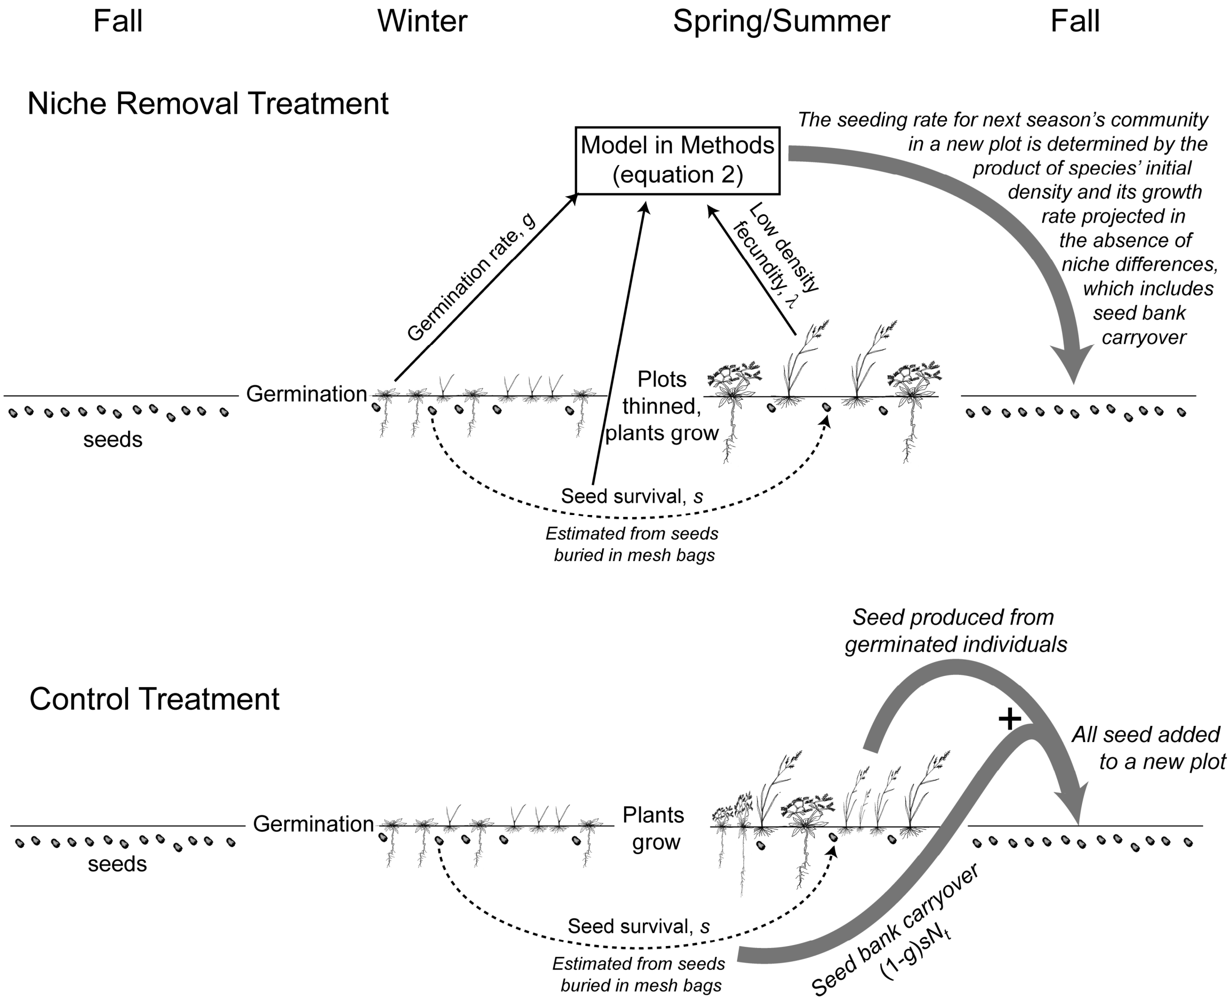

Supplementary Figure 1. **Schematic of the experimental approach**. We compared community dynamics in the presence and absence of niche differences, the results of which are shown in Fig. 3. In both treatments the communities began with 15 g of seed per m^2, evenly divided among the ten competitors. In each niche removal plot, we measured the demographic rates necessary for quantifying population growth rates without niche differences. To determine next year's seeding rate, we multiplied each species' predicted growth rate (which includes seed bank carryover) by its initial seed number, and normalized this product so total community seed mass was 15 g/m^2. The resulting community was then sown into a new plot adjacent to the previous year's community, and the process was repeated. In the control treatments, we measured total seed production by harvesting all seeds off the plots. We estimated the seed bank carryover for each species from the fraction of seeds that did not germinate, seed survival, and the initial sowing density. We summed seed bank carryover with seed production normalized so total community seed mass was 15 g/m^2. The resulting community was then sown into a new plot adjacent to the previous year's community, and the process was repeated. See the text and methods summary for more on the manipulation.

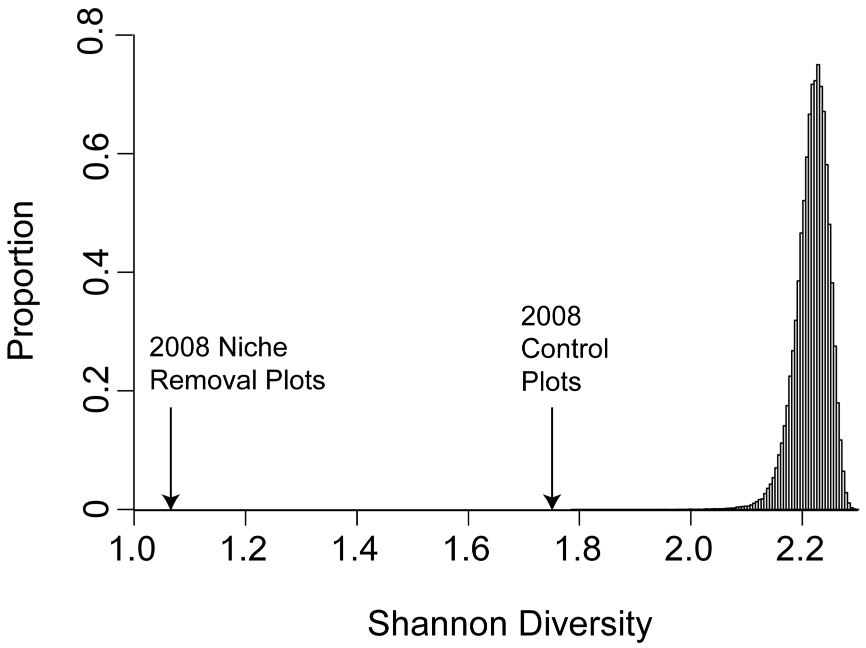

Supplementary Figure 2. Neutral model predictions. We compared the Shannon diversity observed in control and niche removal communities with the predictions of a neutral model. The neutral model tracks the fate of ten identical competitors with demographic stochasticity in all life stage transitions. The number of individuals in the simulations was equivalent to that in the experimental communities (see Supplementary methods for details). Shown is the distribution of Shannon diversity based on 100,000 simulations of the neutral model for two generations of community change, matching the duration of the experiment. The probability of Shannon diversity dropping to levels seen in either experimental treatment was less than 0.0001. Similarly, the probability that the neutral processes in our model could drive a 0.67 difference in Shannon diversity between the ten plots in each treatment was also less than 0.0001.

Supplementary Methods

Model selection

Equation 1 of the Methods is among the most commonly used models for annual plant populations with a seed bank. To ensure that it provided a reasonable description of dynamics in our experimental communities, we used maximum likelihood to compare it to other candidate models describing how seed production changes with germinant density. We used Law and Watkinson's[1] analysis of annual plant competition to provide the candidate models, which included several Beverton-Holt, Ricker, and Lotka-Volterra forms (Supplementary Table 2). For some models, we rescaled the competition coefficients to match their formulation in our study.

We compared the ability of the models to describe how the number of seeds produced by an individual germinant changes with the density of neighboring plants in our experimental communities. In late winter 2008, for each target species, we located nearly monospecific neighborhoods (15 cm diameter) of each competitor species occurring at a range of germinant densities (facilitated by the different experimental treatments). Later in the year, we measured the seeds produced by the target individuals and corrected those values for seed viability. We then used maximum likelihood with a lognormal distribution to fit each of the eight models to the observed data for each pair-wise species interaction (how the fecundity of species *i* changes with increasing germinant density for species *j*). Finally, we used the likelihoods generated by the model fitting exercises to calculate AIC's summed across all pair-wise interactions. Results (Supplementary Table 2) showed that of the eight candidate models, the formulation used in our study (equation 1, Methods) best described how seed production changes with competitor density.

Theoretical projections of species' per capita growth rates without niche differences

Theoretical work by Chesson (2000) shows how the per capita growth rate of a species when it is rare and its competitor(s) common can be decomposed into a scaled sum of its niche difference (stabilization term) and fitness difference from competitors. The niche difference term can be eliminated by forcing species to limit themselves and their competitors equally, and the resulting growth rates reflect the average fitness difference. Without niche differences, the species with the highest average fitness displaces all other competitors[2], and species per capita growth rates are independent of their commonness and rarity, as long as the community is filled with individuals[3, 4]. By filled, we mean that species *i* is rare and *j* is at its carrying capacity, species *j* is rare and *i* is at its carrying capacity, or somewhere in between.

To solve for the population growth rate of species *i* in the absence of niche differences, we first ensure that its competitor *j* has equilibrated in response to the density of species *i*. For competitors *i* and *j* with dynamics described by equation 1 (Methods), we can solve for the equilibrium density of species *j*, N_j^*, in a system with any fixed density of species *i*, by setting the multiplicative rate of increase of species *j* equal to 1.

$$1 = s_j\left(1 - g_j\right) + \frac{\lambda_j g_j}{1 + \alpha_{jj} g_j N_j^* + \alpha_{ji} g_i N_i} \qquad (3)$$

Algebra yields:

$$N_j^* = \frac{1}{\alpha_{jj} g_j}\left[\frac{\lambda_j g_j}{1 - s_j(1 - g_j)} - 1 - \alpha_{ji} g_i N_i\right] \tag{4}$$

To solve for the growth rate of species i in the absence of niche differences we assume that species limit themselves and their competitors equally, meaning we set $\alpha_{jj} = \alpha_{ij}$ and $\alpha_{ii} = \alpha_{ji}$. Using this condition, substituting expression 4 for $N_{j,t}$ in equation 1 (Methods), and simplifying yields equation 2 of the Methods.

We evaluate the accuracy of our approach for projecting population and community dynamics without niche differences by comparing results to an alternative method only available after all years of data have been collected. The latter involves setting all competition coefficients equal to one another, and then directly simulating the multi-species system described by equation 1 (Methods) but expanded for ten competitors. The simulation must be parameterized with the temporal average of species' demographic rates, because it otherwise includes the stabilizing effects of temporal niches- the storage effect and relative nonlinearity[2]. Thus, for each species' demographic parameter (g_i, s_i, or λ_i), we used the geometric mean of its 2007 and 2008 value. We set all of the competition coefficients (α's) equal to 1, and initiated the simulation with each species at one tenth of the average carrying capacity of the ten competitors. Note that because this alternative approach requires temporal averages of the demographic rates, it cannot be used to project community dynamics on a year-specific basis, as required for our experiment in Fig. 3 (Supplementary Fig. 1). Still, a retrospective analysis can be used to validate our approach. To ease comparison of the two methods we repeated the simulation in Fig. 2b with temporally averaged parameters.

After two generations of community change, the time scale of our experiment, the two approaches produced similar relative abundances for the ten competitors. The method used in the text (equation 2) accurately predicted the ten species' relative abundance determined via the alternative method ($R^2 = 0.97$, $P < 0.001$, regression slope anchored at origin = 0.86). Over longer time scales, the approach in the text and methods yielded 99.9% *Salvia* in 15 years, somewhat sooner than the alternative method, which took 29 years. Nonetheless, exclusion within 30 years is still highly inconsistent with coexistence via competitive equivalence.

A neutral model for the experimental communities

To evaluate the possibility that demographic stochasticity alone could explain the observed decline in Shannon diversity in control and niche removal plots (Fig. 3), we developed a neutral simulation model for the system. The model predicts the dynamics of ten competitors equivalent in their average vital rates, but these vital rates vary between individuals.

To parameterize this model, we quantified variation between individuals in seed production with 25 isolated plants per species distributed over five replicate plots in 2006-2007. For each individual, we measured the number of seeds produced per surviving germinant, and corrected that number for seed viability. Individual seed production proved log-normally distributed (the mean and standard deviation of $\log_e$ (seed production) were 6.20 and 0.76 for *Chaenactis*, 6.75 and 1.22 for *Chorizanthe*, 6.62 and 0.97 for *Lasthenia*, 4.45 and 1.35 for *Lotus*, 3.91 and 1.41 for *Micropus*, 10.20 and 0.67 for *Navarretia*, 5.97 and 2.10 for *Plantago*, 7.35 and 1.32 for *Salvia*, 3.67 and 1.76 for *Trifolium*, 5.53 and 1.41 for *Vulpia*). Our neutral simulation thus involved drawing seed production values for individuals from a log normal distribution with a mean and standard deviation matching the ten species average (6.07 and 1.30).

The simulation began with seed numbers for ten species matching those initially sown into our experimental plots (ranging between 150 to 4500 per species depending on species' seed mass). For each species, the number of germinating seeds was drawn from a binomial distribution where the number of trials (n) was the number of seeds and the probability (p) was the ten species average germination rate in our experiments (0.234). The number of seeds carrying over to the following fall in the seed bank was similarly drawn from a binomial distribution where n was the number of ungerminated seeds and p was the average buried seed survival rate in our experiments (0.379).

The contribution of seed production to next year's population was determined as follows: The number of germinants surviving to make at least one seed was drawn from a binomial distribution where n was the number of germinated seeds and p was the average germinant survival to flowering in our experiments (0.778). For each species, all of their surviving germinants were assigned a seed production value based on a lognormal distribution with mean and standard deviation equivalent to the ten species average of these values. As in our experiments, we assumed that the communities could only contain 15 g of seed/m^2. Thus the summed seed production across the ten species was scaled so that when added to seed bank carryover, communities contained the appropriate seed mass (7.5 g of seed for our 0.5 m^2 communities). Our neutral model thus assumes that plants compete, but all species have their seed production equivalently reduced by competition (the scaling term is equivalent across the ten competitors). The process was then repeated for a second year, matching the duration of our experiments.

To evaluate the probability that communities "drift" to Shannon diversity values seen in the control and niche removal plots of our experiment, we ran the above simulation 100,000 times. We then calculated the proportion of simulations in which species abundances drifted sufficiently to generate Shannon diversity values less than or equal to those seen in our two treatments. From these 100,000 simulations, we also calculated the probability that the composition of ten communities in each of two treatments would drift sufficiently to generate a 0.67 unit difference in Shannon diversity (that observed in Fig. 3).

Supplementary References

1. Law, R. & Watkinson, A. R. Response-surface analysis of two-species competition: an experiment on *Phleum arenarium* and *Vulpia fasciculata*. *J. Ecol.* **75**, 871-886 (1987).
2. Chesson, P. Mechanisms of maintenance of species diversity. *Annu. Rev. Ecol. Syst.* **31**, 343-366 (2000).
3. Adler, P. B., HilleRisLambers, J. & Levine, J. M. A niche for neutrality. *Ecol. Lett.* **10**, 95-104 (2007).
4. Levine, J. M., Adler, P. B. & HilleRisLambers, J. On testing the role of niche differences in stabilizing coexistence. *Funct. Ecol.* **22**, 934-936 (2008).