Supporting Information

A minimal model of a forest after neutral theory

Matthew D. Schwartz
Department of Physics, Harvard University

James P. O'Dwyer
Department of Plant Biology, University of Illinois at Urbana-Champaign

S1 Species and data

Purpose: the species we model and those we set aside, the stage classes and stem populations, the regional plots used for the immigration weights, and whether the results depend on the eight published stages.

Jops et al. (2025) published projection matrices for 90 species at Barro Colorado Island (BCI) (Jops and O'Dwyer, 2024). The BCI censuses of 1982–2015 (Condit et al., 2019) list 255 species among the recruits and deaths of trees of at least 10 cm dbh. We model the 84 species on both lists; by the plot's species codes they are acaldi, alchco, alibed, alsebl, annoac, apeime, aspicr, beilpe, brosal, calolo, cappfr, caseac, cassel, chr1ec, chr2ar, coccma, cordbi, cordla, cou2cu, crotbi, cupasy, des2pa, drypst, ery2ma, eugeco, eugega, eugene, eugeoe, faraoc, gar2in, gar2ma, guargu, guatdu, gustsu, hassfl, heisco, hirttr, hybapr, ingago, ingama, ingaqu, ingas1, ingaum, laciag, laetth, licapl, loncla, maquco, micoaf, micoar, mourmy, ocotce, ocotwh, pentma, picrla, poular, poutre, pri2co, protco, protpa, protte, psycho, pterro, quaras, randar, rinosy, simaam, sipapa, soroaf, stylst, swars1, swars2, tab1ro, tab2ar, tachve, talipr, tet2pa, tri2pa, tri2tu, tripcu, tropra, unonpi, virose, xyl1ma. The six species with a published matrix but no recruit or death at that size are anaxpa, herrpu, micone, ouralu, paligu, psycma (*Anaxagorea panamensis*, *Herrania purpurea*, *Miconia nervosa*, *Ouratea lucens*, *Palicourea guianensis* and *Psychotria marginata*), all understory shrubs and treelets. Each has hundreds of censused stems in every census, never fewer than 251, but none had a live stem of at least 10 cm dbh at any of those censuses, their largest stems measuring 5 to 9.9 cm. We therefore do not use their matrices, which were fitted on smaller stems. Table S1 gives the lower diameter edges of the stage classes, among them the large-tree classes (stems of at least 79.4 mm dbh), and the stem populations used in the main text and in this supplement.

The immigration weights come from the network of census plots across Panama (Condit et al., 2022, 2016). We exclude BCI itself and the six plots on or beside the island, and take a species' regional share p_s as its share of the reproductive-size stems in the remaining plots. The nine modeled species absent from those plots enter the pool at the smallest share any sampled species receives, the floor share. Section S2 gives the stationary composition with them at zero weight instead.

To test whether the results depend on the eight published stages, we rebuild the diameter classes from the census stem measurements under eight alternative structures and re-derive three results under each (Table S2). The three results are the out-of-sample decadal forecast of Section S5, the stationary composition under the regional pool, and the ordering across species of the local gain $G_s \propto \Omega_s/\mu_s$, the large trees standing per arriving seed. Most of the structures change only the

class k	2	3	4	5	6	7	8
lower edge (mm)	10	20.0	39.8	79.4	158.5	316.2	631

stems	population
201,000	the censused stand: classes 2–8 (stems of at least 10 mm, the 1 cm census threshold), the 84 modeled species, mean over the censuses of 1982–2010 at which stage vectors are built; the model's capacity, the fixed total of censused stems, is this mean unrounded
21,700	large trees: classes 5–8 (stems of at least 79.4 mm), the 84 species, mean over the same censuses
21,000–22,200	large trees, the range of the per-census totals over the same censuses
14,900–35,300	inferred flow of stems from class 1 into class 2 per census interval over 1982–2010 (mean 22,300)

Table S1: Stage classes and stem populations at Barro Colorado. Top: lower diameter edges of the model's stage classes. Class 1 holds stems below 10 mm, the seedling pool, whose rates Jops et al. (2025) fitted from the annual seedling census; classes 2 through 8, the censused classes, are log-spaced with lower edges $10 \cdot 10^{0.3(k-2)}$ mm and class 8 unbounded above; classes 5 through 8, stems of at least 79.4 mm dbh, are the large trees whose abundances we compare with the model's predictions. Bottom: the stem populations used in the main text and in this supplement, rounded, each with its threshold, species set and censuses.

class boundaries above the seedling pool: fewer and wider log-spaced classes, finer ones at half the published width, ten log-spaced classes from 10 mm to the largest stem, equal-width diameter classes with and without an edge at 79.4 mm, and a finer division of the smallest censused stems. The other two are integral projection models on a continuous diameter axis, one with parametric vital rates (logistic survival and Gaussian growth, both linear in log diameter) and one nonparametric, drawing a stem's fates from those observed for the species' own stems of similar diameter.

In every structure we keep the published seedling pool and its rates, because the five-year census does not measure stems below 10 mm. We rebuild the fecundity schedule with the published size dependence and rescale it to restore each species' long-run growth rate to one. Where a species has too few stems in a class we use the rate pooled over all species, and we hold the immigration weights, arrival rate and capacity fixed. We do not re-derive the results that involve the shared year effect. The transition through which it reaches the large-tree classes, sapling promotion past 79.4 mm dbh, is nevertheless common to all the structures. Each keeps that diameter as a class edge or approximates it by its nearest edge (Table S2, note), and the two continuum models represent promotion past it by the kernel mass that crosses it.

We count the forecast as changed if its rank correlation differs from the value under the published classes by more than 0.05 at both horizons (one and two censuses ahead). The composition under the regional pool counts as changed if its rank correlation with the census differs by more than 0.05 or its identity-line R^2 by more than 0.10, and the gain ordering if its rank correlation with the ordering under the published classes falls below 0.90. By these criteria the gain ordering is unchanged under every structure (lowest rank correlation 0.97), and only the parametric continuum changes the forecast, lowering its rank correlation to 0.628 one census ahead and 0.622 two ahead. The two structures without an edge at 79.4 mm, whose large-tree threshold lies higher (Table S2, note), lower it by slightly more than 0.05 one census ahead only. The largest rise under any structure is +0.013 (the nonparametric continuum, one census ahead), and in every structure the forecast's rank correlation remains above that of each species' own trend.

Under the regional pool, the rank correlation of stationary with observed abundance changes by less than 0.05 in every structure. In contrast, the identity-line R^2 rises by more than 0.10 under the two equal-width structures and the ten log-spaced classes, and falls by more than 0.10 under the parametric continuum, where the forecast's rank correlation is also lowest. We attribute the rise under the equal-width structures to the narrower spread of predicted abundances that a wide first censused class produces. Because the same coarsening also lowers the forecast's rank correlation, the higher R^2 is unlikely to mean that these structures describe the forest better. Under every structure, as under the published one, the stationary order of species broadly matches the observed order, while the predicted absolute abundances remain far from those observed (negative R^2 in every row of Table S2).

		forecast ρ		own trend ρ		regional pool		gain
structure	stages	1	2	1	2	ρ	R^2	order
published, seven log classes	8	0.712	0.683	0.541	0.506	0.667	−0.737	1.000
fewer, four classes	5	0.711	0.674	0.541	0.506	0.682	−0.652	0.995
finer, thirteen log classes	14	0.708	0.684	0.541	0.506	0.663	−0.785	0.994
ten log classes†	11	0.656	0.654	0.501	0.485	0.672	−0.432	0.993
equal width, 69 mm	10	0.674	0.652	0.541	0.506	0.704	−0.571	0.973
equal width, 50 mm†	10	0.660	0.648	0.505	0.481	0.702	−0.243	0.977
finer sapling classes	10	0.713	0.683	0.541	0.506	0.662	−0.816	0.996
continuum, parametric rates	121	0.628	0.622	0.541	0.506	0.646	−0.969	0.978
continuum, nonparametric rates	121	0.725	0.684	0.541	0.506	0.657	−0.756	0.986

Table S2: Three results under each stage structure. The first row is the published structure rebuilt by the same procedure as the alternatives; its forecast rank correlations round to the values reported in the main text. Stages: the number of stages including the seedling pool. Forecast: rank correlation of predicted with observed change in large trees one and two censuses ahead over the 252 species-start pairs (one species at one starting census) of the out-of-sample test (rates from 1982 to 1995 applied at 1995, 2000 and 2005; Section S5), with each species' own early trend, extrapolated, as the comparator. Regional pool: rank correlation and identity-line R^2 of stationary against census-mean log abundance under immigration weighted by regional share, from the closed form of Section S2, with the nine species the regional plots miss at zero weight; entering them at the floor share instead changes the R^2 values by a few hundredths and leaves every rank correlation unchanged. Gain order: rank correlation across species of Ω_s/μ_s under the structure with its value under the published classes. †No edge at 79.4 mm: the large-tree threshold moves to the nearest edge, 90.6 mm for the ten log classes and 100.0 mm for the 50-mm classes, so the large-tree population in these two rows differs from that of the other rows.

S2 The stationary state and the simulations

Purpose: the factorization of the closed-form stationary state and its consequences; the stochastic simulations and their agreement with it; and how the year effect, the treatment of species absent from the regional plots and the arrival rate bear on the stationary state.

Without the year effect, and with every stem-by-stem draw replaced by its expectation (the mean-field approximation of Material and methods 2.2), the stationary stage vector $\bar{n}_s$ of species s is the fixed point of the resulting deterministic map, written in the main text as a matrix inverse. Because the recruitment lottery is nonlinear this fixed point is not the exact expectation of the stochastic model; how closely the stochastic simulations reproduce it is stated below. In that

notation, U_s and F_s are the published survival-and-growth and fecundity matrices of species s (Jops et al., 2025), F_s being nonzero only in its first row, and e_1 is the unit vector on the seedling class (class 1). The expected number of immigrant seeds of species s per five-year interval is Mw_s, with M the arrival rate and w_s the species' share of it. We write d for the expected fraction of candidate recruits not admitted to the censused stand in a step, the same for every species. The matrix $U_{d,s}$ is U_s with its class-1 promotion entries multiplied by $1-d$ and the removed probability returned to the class-1 diagonal, so that a seedling not admitted stays a seedling with its mortality unchanged. The stationary vector then satisfies $\bar{n}_s = (U_{d,s} + F_s)\,\bar{n}_s + Mw_s e_1$. The capacity fixes d through the condition that $\bar{n}_s$, summed over species and over the censused classes 2 to 8, equal the observed number of censused stems. At Barro Colorado the resulting d is 0.153 under the regional pool and 0.026 under the compensated scheme.

Because every published matrix $U_s + F_s$ has a long-run growth rate of one, the matrix inverse that gives $\bar{n}_s$ factorizes. We write ϕ_s for the stable stage vector of $U_s + F_s$ (distinct from the autocorrelation ϕ of the year effect) and μ_s for a seedling's probability of dying in a step. Let $\Omega_s = \sum_{k=5}^{8} \phi_{s,k}/\phi_{s,1}$ be the number of large trees (stems ≥ 79.4 mm dbh, classes 5 to 8 of the published matrices) per seedling in that stable structure. Then

$$\bar{n}_s = M\, w_s \left[\frac{e_1}{\mu_s} + \frac{1-d}{d\,\mu_s} \frac{\phi_s}{\phi_{s,1}} \right], \tag{S1}$$

as substitution into the fixed-point equation above confirms, using $(U_s + F_s)\,\phi_s = \phi_s$ and the facts that $U_{d,s}$ differs from U_s only in its first column and that F_s has no entry in that column (seedlings set no seed). Above the seedling class the stationary stage vector of every species is therefore proportional to its own stable structure ϕ_s, whatever the abundances of the other species. The immigration weight, the arrival rate, the recruitment lottery (through d) and any rescaling of fecundity set only the constant of proportionality.

Summing Eq. (S1) over classes 5 to 8 gives the expected large-tree abundance $\bar{c}_s = M\, w_s\, G_s$, with $G_s = (1-d)\,\Omega_s/(d\,\mu_s)$ the local gain of the main text (large trees standing per arriving seed). The large-tree composition, $w_s\Omega_s/\mu_s$ normalized over species, contains neither d nor M: under the regional pool ($w_s = p_s$, the regional relative abundance) it is $p_s\Omega_s/\mu_s$ normalized, and under the compensated scheme ($w_s \propto p_s\,\mu_s/\Omega_s$, equivalently p_s/G_s) it is the regional composition p_s itself. The arrival rate cancels from every censused abundance as well, because the capacity condition constrains only the second term of Eq. (S1), summed over species and over the censused classes, and thereby fixes the product $M(1-d)/d$. A change in M is therefore absorbed by a change in d and, through the first term, alters only the size of the uncapped seedling pool. In class 1 the two terms sum to $1/(d\,\mu_s)$ standing seedlings per arriving seed, the quantity used in Results 3.4. Of these the first term, the immigrant seeds themselves, is the fraction d and the second, seedlings descended from the standing trees, the fraction $1-d$; since the capacity fixes $M(1-d)/d$, the stationary seedling pool is proportional to $M/d = M + M(1-d)/d$ and responds to the arrival rate only through its first part. Halving M lowers the predicted seedling density by the fraction $d/2$ ($d = 0.153$ under the regional pool, 0.026 under the compensated scheme), while every censused abundance is unchanged, and no value of M lowers it by more than the fraction d. The seedling density of Results 3.4 is in this sense a prediction of the matrices and the capacity, set mainly by the seed production the standing trees must have at growth rate one, and only its last part scales with the nominal seed production B that enters $M = \hat{m}B$.

In the stochastic simulations we follow every stem individually through the dynamics of Material and methods 2.1 in five-year steps. At each step we assign every stem a fate (death, stasis, promotion or regression) by one multinomial draw with its species' published probabilities, and add Poisson seed production and Poisson immigration with expectation Mw_s. The places opened by deaths

among censused stems are then filled at random, without regard to species, from the seedlings drawn for promotion, and seedlings not admitted remain in class 1. When the year effect is included, the sapling-promotion entries of every U_s are multiplied by that step's value of the common factor g_t before the draws.

Every simulation starts far from the stationary state, in one of two families of initial conditions. In both, we place each species at its stable stage distribution ϕ_s and scale the whole community to put the censused stand at 2.7 to 3.9 times its capacity in the first family and well below it in the second. The stand declines to the capacity through mortality or fills to it through recruitment, and we treat a state as stationary only where the two families agree. Under the regional pool we ran 12,000 steps in 36 replicates from each family, 72 simulations in all. We divided each into 30 windows of 400 steps and compared species' mean abundances with the census over the last five of them. The two families agreed from window 5 onward, and at stationarity the number of censused stems equaled the capacity in every simulation.

Evaluated with the immigration weights, arrival rate and capacity of each set of simulations run without the year effect, Eq. (S1) reproduces the large-tree total of every simulated stationary state to within 0.5% and its species abundances at a rank correlation of 0.997 or better; species by species the simulated long-run mean differs from it by 0.6% for the median species under the regional pool (2.1% under a weighting by demographic variance alone, the largest of the weightings we ran). The one route by which a species' demographic variance (Jops et al., 2025) could act on its abundance other than through the mean field is this residual, and it shows no trend: regressed on the logarithm of demographic variance, the logarithm of the ratio of simulated to closed-form abundance changes by at most 0.03 ($\log_{10}$ units) from the lowest to the highest demographic variance among the species, a range of more than three orders of magnitude, under every weighting we ran. The advantage of low-variance life histories is therefore the mean-field one contained in the local gain, which across the 79 operators that reach the large-tree classes falls with demographic variance at a log-log slope of -0.88 (rank correlation -0.78), and demographic noise adds nothing detectable to it at these population sizes. Demographic noise, the randomness of the recruitment lottery and the initial conditions therefore matter here only for judging when a simulation has become stationary. The stationary abundances given in Results 3.2 are those of Eq. (S1) without the year effect.

When the year effect is included, acting on the sapling promotions with the amplitude and memory estimated from the transition data (Material and methods 2.5 and Section S3), the stationary composition of the simulations remains that of Eq. (S1). The rank correlation of the stationary species abundances with the census changes by $+0.003$ under the regional pool and by -0.025 under the compensated scheme, at most 0.03 in size. The leading species' share of the large-tree stand changes by less than one percentage point (0.9 at most). This invariance is a numerical property of the model at the estimated year effect, and we have no analytical argument that it holds in general. The year effect instead changes the size of the fluctuations about the stationary state and, under the compensated scheme, the mean large-tree total (Section S8).

In the closed-form values of the main text, and in this supplement except where stated, the nine species absent from the regional plots take the smallest share w_s that any sampled species receives (the floor share, Material and methods 2.1). In the stationary simulations under the regional pool these species had zero weight instead. With zero weight in place of the floor share in Eq. (S1) itself, the stationary rank correlation with the census is unchanged under both immigration schemes. The number of species predicted below one large tree rises from 16 to 17 under the regional pool and from 5 to 11 under the compensated scheme. The identity-line R^2 on log abundance falls from -0.693 to -0.737 and from $+0.495$ to $+0.396$, because a species with no arrivals has a predicted abundance of zero where the census has a few stems.

We use one arrival rate under both immigration schemes, $M = 73{,}729$ immigrant seeds per five-

year interval, the product of Hubbell's immigration probability fitted to the plot's own censuses ($\hat{m} = 0.026$) and a nominal whole-stand seed production B of 2.80 million seeds per interval (Material and methods 2.3). By Eq. (S1) the composition of the censused stand does not depend on this rate. When we varied it in simulations under the regional pool over a seventeen-fold range, from 1.3% to 18.7% of the stationary input to the seedling class, the rank correlation of the stationary abundances with the census changed by no more than 0.005. Under our model the rate therefore cannot be estimated from censused abundances and would have to be measured directly. The censuses do place a lower bound on it: 30 of the 16,982 recruits over the series belonged to species not previously recorded on the plot, an immigrant fraction of at least 0.0018.

S3 The year effect estimated from the transition data

Purpose: the statistics behind Fig. 2 of the main text; the penalized comparison of year-effect structures, mortality included, and simulations with each species' own mortality variation added; how the amplitude of the adopted year effect is estimated and why its memory is assumed; simulations with a shared factor placed on each life-cycle transition in turn; and exploratory comparisons of the species' responses to the shared effect with life history, light response and climate.

The data are the plot's censuses of 1982–2015 taken as seven intervals, the first of three years and the rest of five. For each of the 84 species and each interval we count the stems of each diameter class promoted to a higher class among those alive in the censused classes at the second census. We also count the deaths in classes 2 to 4 and in classes 5 to 8, and the stems that crossed the 10-mm line per stem standing at the first census. These counts give nine rates per species: three sapling promotions (out of classes 2, 3 and 4, from 10 mm dbh to the large-tree threshold of 79.4 mm), three tree promotions (out of classes 5, 6 and 7), two mortalities and recruitment. A species' deviation in an interval is the complementary log-log of its promoted or dead fraction, less the logarithm of the interval length in five-year units, centered on the species' seven-interval mean; for recruitment it is the centered logarithm of the rate per stem per year. An interval effect on this scale is a log hazard ratio, close to the logarithm of the factor g_t of equation (2) of the main text.

The component common to all species in a rate is the set of centered interval effects of a binomial regression (Poisson for recruitment) with species and interval terms, fitted to every species with events. Table S4 gives the standard deviation of that component for each rate beside its sampling floor (the same statistic simulated with no year effect). A species is included in Fig. 2 of the main text and in Table S3 for a rate when its mean rate implies at least twenty events per interval and it has stems at risk in every interval. Sapling growth varies synchronously across species: 99% of the 175 species-by-promotion series correlate positively with the component common to the other species, 69% above the seven-point null, and the median species follows it at close to full amplitude. Mortality in classes 2 to 4 varies between intervals within a species as much as sapling growth does but is not synchronous across species (median r +0.44, 27% above the null), and its common component is correspondingly small.

To decide which structure of year effects these data support, we fit every species-by-interval count of every rate with stems at risk, 4,228 observations over nine rates and seven intervals (4.0 million stem-intervals), the same for every model. Promotion and death counts are fitted by binomial likelihood with a complementary log-log link and the log interval length as offset, and recruitment counts by Poisson likelihood. Every model has one intercept for each species and rate, and the models differ only in the year-effect terms listed in Table S5; each is also refitted on the six intervals of 1985–2015, because the 1982 census measured small stems by a different method. At this sample size the Bayesian information criterion favors a year effect of some form in each

group of rates (sapling promotions, tree promotions, mortality and recruitment). It rejects only two elaborations, free interval effects for the tree promotions in place of one loading and free species-by-interval effects in place of a variance component. The composite model M6 has the lowest BIC under either choice of sample size (counts or stem-intervals) and on both the seven and the six intervals. One shared effect on sapling growth is therefore not a statistically complete description: tree growth, mortality and recruitment all vary between intervals, species deviate from the common interval effects by detectable amounts, and the three sapling promotions differ in amplitude (model M1b) on seven intervals, though not on 1985–2015 ($P = 0.46$).

Judged instead by the amplitude of each term (Material and methods 2.5 of the main text), the structure of the year-to-year variation is simple. The shared sapling effect has standard deviation 0.20 over the seven intervals. The tree promotions load $\lambda_{\text{late}} = 0.86$ on that effect (90% profile CI 0.75–0.97) and, given interval effects of their own, reproduce it (correlation 0.95), so that growth follows the same interval effects at every size, although this conclusion rests mostly on the promotion out of class 5. Of the systematic between-interval variance of the sapling promotion rates, 71% is common to all species, 10% is species-by-interval deviation coherent across a species' sapling stages ($\sigma_b = 0.075$), and 20% is count-level overdispersion. Mortality loads -0.06 on the shared growth effect in classes 2 to 4 and $+0.39$ in classes 5 to 8, under half the growth amplitude, and its own interval effects (standard deviations 0.09 and 0.16) are only weakly correlated with those of growth. On the basis of these amplitudes we adopt one factor g_t per interval, common to species and acting on the three sapling promotions with one amplitude, with no factor on mortality and no species-specific response.

The opposite assumption, that every species responds to the environment in its own way, corresponds to model M4F, whose 498 parameters beyond the six of the shared effect raise the log-likelihood by 909, or 3.6 units of deviance per parameter, compared with 932 per parameter for the six common effects. Akaike's criterion (two units per parameter) retains the species responses, the Bayesian criterion (8.35 per parameter here) rejects them, and one number for their spread, σ_b, is supported by both. The transition data thus contain species-specific growth responses, but these amount to a tenth of the systematic between-interval variance, and we accordingly use the shared effect alone. For the two mortality rates we have not made the corresponding comparison of species-by-interval terms (the counterparts of M4F and M4R); we test the consequence of the species-level mortality variation of Table S3 by simulation instead.

A mortality term that varies between intervals independently in each species has no persistent direction and therefore, apart from a slight rise in mean survival (a mean-one factor on the hazard lowers the five-year death probability, by 0.4% of its value at the measured amplitudes), leaves each species' mean matrix unchanged, together with the forecast and the stationary abundance that follow from it. It adds instead to the size of each species' census-to-census changes, the quantity in which the model falls short of the census for species of intermediate abundance (Results 3.3 of the main text). We add such a term to the stationary simulations of Section S8 at its measured size. At every five-year step we multiply each species' mortality hazard by a lognormal factor of mean one, independent across species and steps, whose log-scale standard deviation is the median species value of Table S3 (0.19 in classes 2 to 4, 0.24 in classes 5 to 8). We simulate the term alone and together with the shared growth factor at $\sigma_g = 0.32$ (Table S25, lower rows).

Simulated alone, the mortality term gives the large-tree total a median five-year change equal to 0.41 of the observed value under the regional pool and 0.48 under the compensated scheme, since variation that is independent across species cancels only partly in the total when a few species hold most of the large trees. By itself the shared growth factor adds 3.6 (regional pool) and 4.3 (compensated scheme) times as much to the total's five-year change as the mortality term does. Added to the shared growth factor, the mortality term alters the total's five-year change under

the regional pool by less than the spread among replicate simulations and lowers it slightly under the compensated scheme. Under that scheme it also restores the standing large-tree total to its value without either term, because a mean-one factor on the mortality hazard raises mean survival. When we rescale each species' baseline hazard in every class so that its expected five-year survival under the term equals the measured value, that restoration disappears (the stand's level is 0.68 of its value without the year effect, against 1.01 before the rescaling and 0.68 with the growth factor alone), and so does the term's slight lowering of the total's five-year change under that scheme; the species-level results that follow are unchanged by the rescaling, the census's mean squared change lying inside the model's range in 6 of the ten abundance bands instead of 7, the species criterion being met in six of the six bands of 16 or more stems under both schemes, and the band values moving by at most 5% under the regional pool.

The mortality term increases the census-to-census fluctuations of every species, most for the abundant species, whose demographic noise is small beside a mortality deviation of a few percent. With both terms the census's mean squared five-year change lies within a factor of two of the model's in the top two of the ten abundance bands, and inside the model's range in more of those bands than with the growth factor alone, under either scheme. Every criterion of Table S25 is then met under both schemes (Section S8), the level criterion under the compensated scheme only through the rise in mean survival described above. For species with 16 to 255 large trees the census's mean squared change still exceeds the model's by about as much as with the growth factor alone, because a mortality deviation of this size is lost in the demographic noise of a species of thirty stems. For these species the species-specific part of the growth variation ($\sigma_b = 0.075$) remains untested as a source of this shortfall.

We estimate the amplitude σ_g of the year effect from the interval effects of the three sapling promotions. A factor that multiplies all three by one g_t implies one common path on the log-hazard scale, which we take to be the mean of the three promotions' centered interval effects, interval by interval (the pooled sapling path, dashed in Fig. 2 of the main text). We summarize that path by its standard deviation over the seven intervals with the regressions' sampling variance removed, $S = 0.222$. The three sets of effects are consistent with a single path: the first principal component of the three-by-seven table of effects accounts for 93% of its variance. The loadings of the three promotions on that component, scaled to average one, are 0.72, 1.24 and 1.03 out of classes 2, 3 and 4, each consistent with one (bootstrap resampling of the census intervals).

Our estimate of σ_g is the value at which a simulated year effect, sampled at the census dates with the census numbers at risk, reproduces S. For each σ_g on a grid we draw 20,000 sequences of g_t with autocorrelation 0.76 per five-year step (the memory assumed below). We apply each sequence to the three promotions at their pooled five-year probabilities, add the census sampling errors to obtain the interval effects that a regression would estimate, and compute the statistic as for the census data. The median simulated statistic equals S at $\sigma_g = 0.318$, and the 90% confidence interval, the set of amplitudes whose central 90% of simulated statistics contain S, is 0.17–0.84. The upper limit is high because seven serially correlated samples of a slowly varying factor often show much less than its full spread. Dropping the first interval, 1982–1985, whose pooled effect is the largest, from both the census statistic and the simulations gives $\sigma_g = 0.27$ (90% CI 0.14–0.71). Every statement about growth made above also holds for 1985–2015 alone (Table S5, last column): the tree promotions then load 0.75 (0.62–0.88), the three sapling loadings on the pooled path are 1.02, 1.00 and 0.98, and the common and species-specific shares of the systematic variance are 80 and 13%.

The memory τ_g of the year effect cannot be estimated from the pooled path, whose lag-one autocorrelation, +0.32, lies inside the range -0.34 to +0.58 that a factor with $\tau_g = 18$ yr gives on seven points, and inside the corresponding range for every correlation time from 5 to 40 years.

Over that range of correlation times the amplitude estimate would change only from 0.23 to 0.43. Differences between species' rates do persist across intervals, and an earlier version of this analysis read τ_g from that persistence; it does not, however, measure τ_g. For each of 255 species and each of the seven intervals we take the net rate per standing tree at which the species gained trees of at least 100 mm dbh by growth past that diameter and lost them by death. The rank correlation of those rates across species between two intervals is 0.33 for consecutive intervals (0.46 over the 84 modeled species) and falls as the intervals lie further apart, to -0.03 at a separation of six; an exponential fitted to the 21 pairs of intervals decays with a time constant of 17.9 yr (90% CI 10.8–33.2 yr from a bootstrap over species). A factor common to all species moves them together and so cannot by itself order them differently from one interval to the next. In 316,800 records of 35 years simulated from the model at the census sample sizes, with the factor's correlation time set to 0, 5, 17.9 and 50 yr and with no factor at all, the same statistic gives median rank correlations between -0.08 and +0.02 at every lag and a fitted time constant of 0.5 to 1.7 yr that does not move with the imposed one; the forecast operators of Section S5, whose growth rates differ among species, give instead a persistence that does not decay (rank correlation 0.51 to 0.57 at every lag). The observed persistence, decaying over two to three decades, therefore belongs to species-specific departures from the mean rates that the model leaves out (the species terms of Table S3), and nothing in the censuses fixes τ_g. We set $\tau_g = 18$ yr, an autocorrelation of $\phi = 0.76$ between consecutive five-year steps, by assumption, and carry correlation times from 5 to 40 yr, over which the amplitude estimate ranges from 0.23 to 0.43. No climate variable we examined has a memory of the assumed length, and the autocorrelation of every candidate series (the covariates of Table S9) decays within 4.4 yr. Recruitment across the 100-mm line varies between censuses with a standard deviation of 18% of its mean rate, nearly all of it common to the species; this value, 0.18, is the measured recruitment variability used in the placement simulations below and in Section S8.

As a test independent of the interval effects estimated above, we also ask on which transitions a factor shared by all species could act and still produce the recruitment variability observed among large trees (Table S6, Fig. S1). We place the year effect on each transition of the life cycle in turn and find the amplitude, up to 1.5 times the mean, at which the modeled between-census variability of recruitment into class 5 (the first large-tree class) reaches the measured 0.18. With the factor on seed production, where environmental variation is usually assumed to act, the variability does not reach the measured value at any amplitude we simulated, and with the factor on establishment or on promotion out of class 2 it likewise levels off below that value. It reaches that value at $\sigma_g = 1.06$ on promotion out of class 3, at 0.30 on promotion into class 5 and at 0.27 on the three sapling promotions together. The last two are the only placements under which the long-run composition remains that of the model without the year effect while the community total varies between censuses as much as the observed total does. With the factor on sapling or adult survival the variability stays below the measured value at every amplitude simulated under the regional pool, while the stand changes in size or in rank order and, when the factor acts on adult survival, its large-tree total fluctuates far more than observed.

Factors placed on the seedling side of the 10-mm line fail to reproduce the measured variability mainly because the variation they impose is attenuated as stems pass slowly through the sapling classes. Each of those classes holds a stem for 13 to 25 years, comparable to or longer than the assumed 18-year memory of the year effect (Fig. S1b). The same attenuation is visible in the census, where recruitment at 10 mm varies by 0.35 of its mean between intervals and the numbers standing in the sapling classes by 0.06 to 0.04. Ingrowth into class 5 nevertheless varies by 0.18, several times more than the class-4 numbers that feed it, which suggests that in the census too the recruitment variability arises at that transition. The common components of Table S4 distinguish between the two placements that reproduce the measured variability without resizing or reordering

the stand. A factor on promotion into class 5 alone implies common components no larger than the sampling floors at the two lower promotions, whereas a factor on all three sapling promotions at $\sigma_g = 0.27$ gives 0.10–0.45 out of class 2 and 0.10–0.40 out of class 3, ranges that contain the census values.

In the placement simulations above, the year effect on seed production is a linear factor clipped at zero, which caps its amplitude near $\sigma = 1.5$. To go beyond that cap we multiply seed production by a lognormal factor of mean one at three larger amplitudes, under the compensated immigration scheme (Table S7). Raising the amplitude mainly thins the stand, whose large-tree total falls to as little as 0.15 of its size without the year effect while the variability of ingrowth into class 5 stays below the measured value. The stand thins because entry into it is limited by the number of places that deaths open. Without the year effect the seedling pool supplies only 3% more candidates than there are places, so that the additional seedlings of a good year mostly cannot enter whereas the missing seedlings of a poor year reduce entry by their full number. Unlike the factor on seed production, a factor on promotion into class 5 at the amplitude found under the regional pool still reproduces the measured recruitment variability under the compensated scheme, lowering the standing total to 0.87 of its value without the year effect (Table S7, last row).

Seed-production noise also fails to supply the variability that the model lacks at intermediate abundance (Section S8). A factor shared by all species increases each species' mean squared change in proportion to the square of its abundance, whereas demographic variance grows in proportion to abundance itself. Relative to demographic noise, the shared factor's contribution is therefore small for all but the abundant species, on any transition and at any amplitude.

The trait and climate comparisons that follow are exploratory and descriptive, and every interval-level statistic in them rests on seven points. For each species we regress its interval deviations on the component common to the other species (the leave-one-out common series of Table S3) in every sapling promotion for which it qualifies, and pool the slopes into one loading per species (76 species). Table S8 gives the rank correlation of the loadings with each of nineteen traits, and also the partial rank correlation given log abundance, because the abundant species dominate the leave-one-out series. The loadings are correlated neither with demographic variance (Jops and O'Dwyer, 2023) nor with stature; the largest correlations, all negative, are with the light-response and fast-growth indices of Rüger et al. (2022) and Wright et al. (2010) and with seedling mortality. If synchronized canopy opening drove the common variation in sapling growth, the light-demanding half of the species should follow that variation more strongly than the shade-tolerant half. However, when the species are split at the median growth response to light, the common series of the two halves are correlated at 0.96, and the amplitude of the light-demanding half's series is 0.88 (90% CI 0.75, 1.01) times that of the shade-tolerant half's. This comparison lends no support to a light-mediated driver, although it cannot exclude one, because the light indices describe contrasts between light environments sustained over census intervals whereas a driver would act through the light of particular years.

Seven intervals can neither establish nor exclude a climate driver, and Table S9 only describes how the pooled sapling path is correlated with published climate series (sources in Section S10). No correlation with the nine covariates is larger than would be expected by chance from so many comparisons on seven points, and the largest, -0.86 with January-to-April radiation, has the sign opposite to that expected under light limitation. The path's rank correlation with the interval midpoint is -0.29 ($P = 0.56$), whereas a steadily changing driver such as rising atmospheric CO_2 or warming would give a value near plus or minus one. Growth measured within census intervals and compared with monthly climate could discriminate among candidate drivers, as could the comparison made here if it were repeated at other census plots in central Panama.

rate	n	size: median s.d. [quartiles]	direction: median r [quartiles]	$r > 0$ (%)	above null (%)	in phase (%)	memory: median ρ_1 [quartiles]
promotion 2 → 3	76	0.19 [0.15, 0.24]	+0.77 [+0.60, +0.85]	100	67	88	+0.02 [−0.13, +0.23]
promotion 3 → 4	62	0.27 [0.22, 0.32]	+0.83 [+0.64, +0.93]	98	73	76	+0.27 [+0.16, +0.36]
promotion 4 → 5	37	0.24 [0.13, 0.32]	+0.77 [+0.64, +0.90]	100	68	73	+0.16 [+0.01, +0.29]
promotion 5 → 6	13	0.12 [0.07, 0.18]	+0.69 [+0.60, +0.83]	100	62	85	+0.04 [−0.07, +0.14]
promotion 6 → 7	3	0.33 [0.28, 0.34]	+0.78 [+0.71, +0.87]	100	67	67	+0.07 [−0.07, +0.13]
mortality, classes 2–4	75	0.19 [0.12, 0.25]	+0.44 [+0.30, +0.68]	95	27	75	+0.20 [−0.11, +0.39]
mortality, classes 5–8	28	0.24 [0.16, 0.31]	+0.60 [+0.49, +0.84]	93	43	68	+0.14 [−0.05, +0.29]
recruitment at 10 mm	82	0.41 [0.33, 0.53]	+0.84 [+0.69, +0.92]	95	77	66	+0.33 [+0.27, +0.40]

Table S3: Census-to-census variation of transition rates for each species, Barro Colorado 1982–2015 (seven intervals; census data). Classes are the eight stages of the published matrices (Jops et al., 2025), class 2 beginning at 10 mm dbh and class 5, the first large-tree class, at 79.4 mm. For each rate, over the n species whose mean rate implies at least twenty events per interval: the median across species of the standard deviation of a species' seven interval deviations on the log-hazard scale, sampling variance removed (size); the median correlation of a species' series with the component common to all other species (direction), the percentage of species positively correlated with it and the percentage whose correlation exceeds the 95th percentile for an unrelated seven-point series (0.67); the percentage best aligned with it at lag zero and not one interval ahead or behind (in phase; an unrelated series is in phase 16–29% of the time); and the median lag-one autocorrelation of the species series (memory), to be compared with the central 90% range of seven points with no memory, −0.62 to +0.36, and of seven points of a process with the year effect's memory, −0.36 to +0.60. No species reaches twenty expected events per interval for the promotion out of class 7 (340 events in all). The weak synchrony of mortality in classes 2 to 4 rests partly on the 1982–1985 interval; on 1985–2015 the species' mortality series are more strongly correlated with one another.

rate	species	sd, all	sd, from 1985	floor	common (%)	$\mathrm{sd}_{\mathrm{common}}$	$\mathrm{sd}_{\mathrm{specific}}$
promotion 4 → 5, per survivor	77	0.232	0.174	0.030	78	0.231	0.122
promotion 4 → 5, per stem at risk	77	0.245	0.178	0.031	77	0.244	0.135
promotion 2 → 3, per survivor	84	0.182	0.186	0.012	86	0.182	0.073
promotion 3 → 4, per survivor	84	0.284	0.176	0.018	81	0.283	0.137
classes 2–4 into class 5	77	0.178	0.162	0.029	36	0.177	0.237
mortality, classes 2–4	84	0.090	0.094	0.011	10	0.089	0.271
mortality, classes 5–8	77	0.158	0.152	0.032	100	0.156	0.000
recruitment at 10 mm, per stem	84	0.364	0.226	0.010	47	0.364	0.388

Table S4: The year component common to the species in each transition rate, Barro Colorado, seven intervals between 1982 and 2015 (census data; these are the bold lines of Fig. 2 of the main text). For each rate: the species included (those with at least five events in at least five intervals); the standard deviation of the centered interval effects of a regression with species and interval terms (binomial with complementary log-log link; Poisson for recruitment), over all seven intervals and over the six from 1985; the 95th percentile of the same statistic with no year effect at the census numbers at risk (the sampling floor, 1,000 simulated tables); the share of the interval-to-interval variance of the species' rates that is common to them; and the common and species-specific standard deviations behind that share. For promotion into class 5, a factor on that promotion alone with $\sigma_g = 0.30$ and $\tau_g = 18$ yr, sampled at the census dates with the census numbers at risk, implies 0.098 to 0.466 (median 0.211; central 90% of 2,000 draws) over seven intervals and 0.088 to 0.454 over six; counting every class-4 stem at risk, 0.098 to 0.468. In each case the census value lies inside the simulated range.

	year-effect terms	df	$\log L$	ΔBIC	ΔBIC (stems)	ΔBIC 1985–2015
M0	none: the mean rates only	0	−31,944	25,397	25,136	11,002
M1	one effect a_t, common to species, on the sapling promotions $2 \to 3$, $3 \to 4$, $4 \to 5$	6	−29,146	19,852	19,633	7,229
M1b	M1 with an amplitude per promotion ($\lambda_{23} = 0.75$, $\lambda_{34} = 1.21$; 1.06, 1.02 on 1985–2015)	8	−29,022	19,619	19,414	7,244
M1c	interval effects free per sapling promotion	18	−28,668	18,996	18,859	7,103
M2	M1 + the same a_t on the tree promotions $5 \to 6$, $6 \to 7$, $7 \to 8$	6	−29,053	19,664	19,445	7,140
M2b	M1 + $\lambda_{\text{late}}\, a_t$ on the tree promotions ($\lambda_{\text{late}} = 0.86$)	7	−29,050	19,668	19,455	7,137
M2c	M1 + interval effects of their own on the tree promotions	12	−29,041	19,692	19,514	7,158
M3	M2 + the same a_t on mortality in classes 2–4 and 5–8	6	−30,172	21,904	21,685	8,976
M3b	M2b + $\lambda\, a_t$ on each mortality (−0.06, +0.39)	9	−29,005	19,594	19,395	7,026
M3c	M2c + interval effects of their own on each mortality	24	−28,400	18,510	18,414	6,084
M1+Mb	M1 + $\lambda\, a_t$ on each mortality, no tree-growth term	8	−29,101	19,777	19,572	7,120
M1+Mc	M1 + own interval effects on each mortality, no tree-growth term	18	−28,505	18,670	18,533	6,155
M5s	M1 + the same a_t on recruitment at 10 mm	6	−21,804	5,167	4,948	1,850
M5	M1 + $\lambda_{\text{rec}}\, a_t$ on recruitment ($\lambda_{\text{rec}} = 1.83$)	7	−20,902	3,372	3,160	1,722
M5c	M1 + own interval effects on recruitment	12	−20,637	2,883	2,704	1,657
M4F	M1 + species × interval effects on the sapling promotions, free	504	−28,237	22,192	25,386	9,320
M4R	M1 + species × interval effects, random ($\sigma_b = 0.099$)	7	−28,803	19,173	18,961	6,791
M4Rk	M1 + count-level random effects ($\sigma_c = 0.128$)	7	−28,468	18,503	18,291	6,790
M4R2	M1 + both ($\sigma_b = 0.075$, $\sigma_c = 0.107$)	8	−28,459	18,495	18,289	6,752
M6−	the supported shared terms together: a_t free per sapling promotion, own effects on recruitment and on each mortality, $\lambda_{\text{late}}\, a_{45,t}$ on the tree promotions ($\lambda_{\text{late}} = 0.79$)	37	−19,414	646	640	356
M6	M6− + the species × interval random term ($\lambda_{\text{late}} = 0.76$, $\sigma_b = 0.099$)	38	−19,087	0	0	0
Mown	every rate its own interval effects (reference)	54	−19,379	718	828	417
Msat	Mown + the free species × interval effects (reference)	540	−18,948	3,915	7,355	2,635

Table S5: Year-effect structures compared on the Barro Colorado transition data. Each model is fitted by maximum likelihood to the same 4,228 species-by-interval-by-rate counts (binomial with complementary log-log link for promotion and death, Poisson for recruitment), each with one intercept per species and rate (604), and the models differ only in the year-effect terms listed; df counts the parameters beyond those intercepts. ΔBIC is the Bayesian information criterion less its lowest value in the column, computed with the number of counts as the sample size and, beside it, with the number of stem-intervals at risk; the last column gives ΔBIC for every model refitted on the six intervals of 1985–2015 (3,624 counts). Estimates in parentheses are for the seven intervals unless marked. M4R2 holds the intercepts and a_t at M4R's values, so its likelihood is a lower bound. The corresponding species-by-interval comparison for the two mortality rates (the counterparts of M4F and M4R) has not been made and has no rows here; the consequence of species-level mortality variation is tested by simulation in Section S8 (Table S25).

transition multiplied by the year effect	σ_g reaching 0.18	CV at σ_g=1.5 if never	stand size (no effect =1)	composition unchanged
seed production	none	0.038	0.88	no
seedling survival (class 1)	none	0.087	0.83	no
establishment into class 2 (crossing 10 mm)	none	0.078	0.77	no
promotion out of class 2	none	0.085	0.91	no
promotion out of class 3	1.06		0.94	no
promotion into class 5 (first large-tree class)	0.30		1.00	yes
the three sapling promotions together (adopted)	0.27		0.98	yes
sapling survival, classes 2–4	none	0.098	1.12	no
adult survival, classes 5–8	none	0.032	0.56	no

Table S6: Placement of a year effect shared by all species on each transition of the life cycle (model). For each transition, under the regional-pool immigration scheme: the amplitude σ_g at which the modeled between-census coefficient of variation (CV) of recruitment into class 5 over seven censuses reaches the value measured at Barro Colorado, 0.18; where no amplitude up to 1.5 reaches it, the CV attained there; the stationary number of censused stems at that amplitude as a fraction of its value without the year effect; and whether the long-run composition remains that of the model without the year effect, by the tolerance on the stationary rank correlation with the census stated in Section S8. The survival placements multiply the mortality hazard. Only the two placements at or just below the large-tree threshold reproduce the measured variability without resizing or reordering the stand. This comparison constrains where a shared factor can act; the amplitude used in the model, $\sigma_g = 0.32$, is estimated from the transition rates (see text) rather than taken from this table.

year effect on	total	CV	bands	5-yr	30-yr
seed production, lognormal, three amplitudes (s.d. 1.3, 2.5, 4.1 × mean)	0.33, 0.19, 0.15	0.08, 0.10, 0.11	5, 6, 7	0.013, 0.016, 0.017	1.08, 1.09, 1.10
seed production, middle amplitude, an independent factor per species	0.48				
promotion into class 5, $\sigma_g = 0.30$	0.87	0.18	4	0.015	1.08

Table S7: Seed-production noise at large amplitude compared with a year effect on promotion into class 5, under the compensated immigration scheme ($w_s \propto p_s/G_s$; model). Rows: the year effect on seed production as a mean-one lognormal factor at three increasing amplitudes, the three values listed in order in each column; the middle amplitude with an independent factor for each species; and the year effect on promotion into class 5 at $\sigma_g = 0.30$. Columns: the standing large-tree total as a fraction of its value without the year effect; the seven-census coefficient of variation of ingrowth into class 5 (measured 0.18); the abundance bands, of ten, whose census value lies inside the model's range over windows (Section S8); the median five-year change of the large-tree total as a fraction of it and its thirty-year range ratio (census 0.0127 and 1.062; bound 1.10).

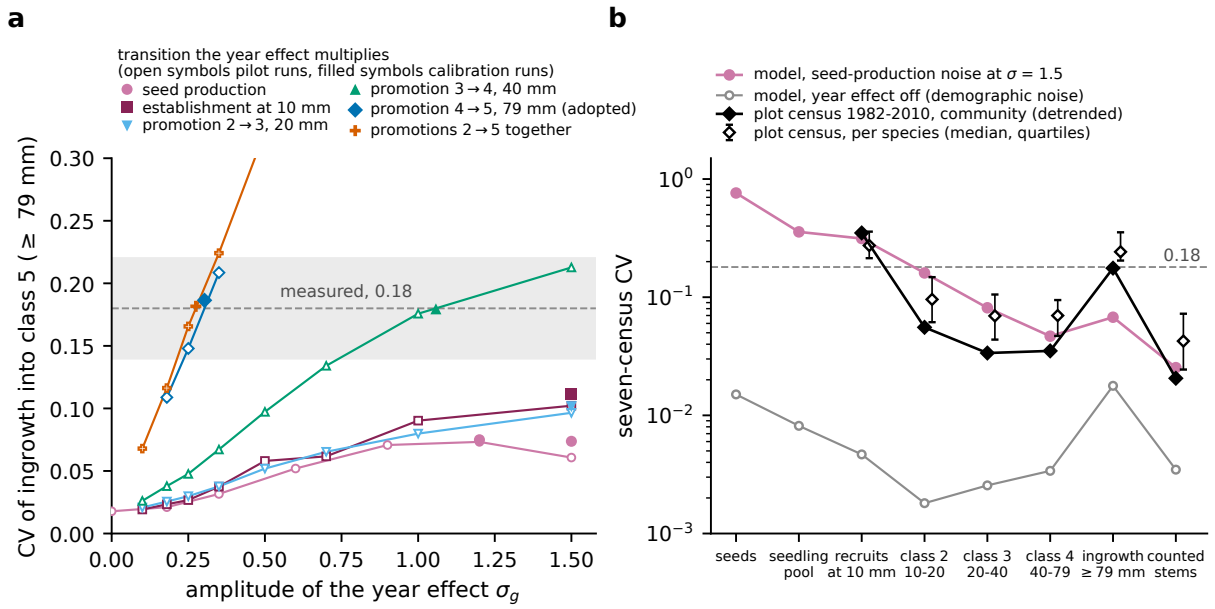

Figure S1: Placement of a year effect shared by all species, and the attenuation of a recruitment signal through the sapling classes. (a) Model: the between-census coefficient of variation of recruitment into class 5, the first large-tree class, over seven censuses, against the amplitude σ_g of the year effect, for each transition the effect multiplies; open symbols are short simulations, filled symbols the longer simulations used to locate each crossing, and the dashed line and band mark the value measured at Barro Colorado (0.18) with its tolerance. Promotion into class 5 and the three sapling promotions together cross the measured value at $\sigma_g = 0.30$ and 0.27; promotion out of class 3 needs 1.06; seed production, establishment at 10 mm and promotion out of class 2 level off below it at every amplitude simulated, up to 1.5. (b) The passage of a recruitment signal up the size classes, in the model and in the census: the seven-census coefficient of variation of seed production, the seedling pool, the stems admitted at 10 mm, the numbers standing in classes 2, 3 and 4, ingrowth into class 5 and the large-tree total, in the model with the year effect on seed production at $\sigma_g = 1.5$ (filled circles) and with no year effect (open circles), and in the plot's censuses of 1982 to 2010 for the whole community (filled diamonds, detrended) and for the median species (open diamonds with quartiles). In both, the variation present at 10 mm has largely disappeared in the sapling classes; in the census, ingrowth at the large-tree threshold nevertheless varies at the measured 0.18, which the model reproduces, without resizing or reordering the stand, only when the effect acts at or just below that transition.

trait	source	n	ρ [90%]	given abundance [90%]
demographic variance	published matrices; Jops and O'Dwyer (2023)	76	−0.10 [−0.30, +0.10]	+0.01 [−0.19, +0.20]
stature (mean height of the six largest stems, else maximum height), m	Wright et al. (2010); Condit et al. (2020)	73	−0.10 [−0.29, +0.11]	−0.12 [−0.30, +0.07]
maximum height, m	Rüger et al. (2022)	76	−0.07 [−0.26, +0.12]	−0.05 [−0.24, +0.13]
wood density, g cm^{-3}	Wright et al. (2010)	72	+0.28 [+0.10, +0.45]	+0.28 [+0.09, +0.44]
seed dry mass, g	Wright et al. (2010)	65	+0.13 [−0.09, +0.34]	+0.05 [−0.18, +0.27]
dry-season moisture association	Condit et al. (2013)	66	+0.02 [−0.20, +0.24]	−0.02 [−0.23, +0.19]
sapling growth potential RGR95SAP	Wright et al. (2010)	72	−0.49 [−0.64, −0.32]	−0.34 [−0.52, −0.13]
mortality of the slowest-growing quarter of saplings MRT25SAP	Wright et al. (2010)	72	−0.30 [−0.47, −0.10]	−0.08 [−0.27, +0.11]
growth response to light (growth in high light over growth in low light)	Rüger et al. (2022)	75	−0.39 [−0.56, −0.21]	−0.19 [−0.39, +0.03]
mortality response to light (high over low)	Rüger et al. (2022)	75	+0.21 [+0.01, +0.39]	+0.06 [−0.14, +0.26]
recruitment response to light (high over low)	Rüger et al. (2022)	75	−0.39 [−0.56, −0.19]	−0.23 [−0.45, −0.01]
growth in shade (mm yr^{-1} of a 10-mm stem in low light)	Rüger et al. (2022)	75	−0.45 [−0.60, −0.28]	−0.32 [−0.49, −0.12]
mortality in shade (yr^{-1} of a 10-mm stem in low light)	Rüger et al. (2022)	75	−0.28 [−0.44, −0.09]	−0.06 [−0.27, +0.15]
annual survival in the large-tree classes	published matrices	71	+0.18 [−0.04, +0.39]	+0.12 [−0.08, +0.31]
probability that a 10-mm recruit reaches the large-tree classes	published matrices	71	−0.08 [−0.27, +0.12]	−0.14 [−0.31, +0.06]
first-class recruits per large tree per year	published matrices	71	−0.08 [−0.28, +0.13]	+0.03 [−0.17, +0.22]
local gain G_s (large trees standing per arriving seed)	published matrices; equation (4) of the main text	71	+0.05 [−0.16, +0.25]	−0.03 [−0.22, +0.15]
large trees per seedling in the stable structure, Ω_s	published matrices	76	−0.08 [−0.28, +0.13]	−0.08 [−0.26, +0.12]
seedling mortality per five-year step, μ_s	published matrices	76	−0.41 [−0.55, −0.25]	−0.32 [−0.47, −0.14]

Table S8: Each species' loading on the common sapling-growth variation compared with nineteen traits, Barro Colorado 1982–2015 (seven intervals); exploratory and descriptive. The loading is the inverse-variance mean, over the sapling promotions for which the species' mean rate implies at least twenty events per interval, of the slope of its interval deviations on the leave-one-out common series (Table S3); 76 species, median loading 0.89, standard deviation 0.30 across species of which 0.19 remains after subtracting sampling variance. n, species with the trait; ρ, Spearman rank correlation with its bootstrap 90% interval (10,000 resamples of species); given abundance, the partial rank correlation given the logarithm of the species' mean stems at risk, with its interval. High and low light are the two reference light conditions of Rüger et al. (2022), whose tables cover 83 of the 84 species; the indices of Wright et al. (2010) cover 76. Of the nineteen intervals in the fourth column, nine exclude zero (1.9 expected by chance); the largest $|\rho|$, 0.49, exceeds the 0.34 reached by the largest of nineteen in 95% of permutations of the loadings across species. Demographic variance and the six matrix functionals are computed from the same published matrices and are correlated among themselves, so their rows are one family and not independent comparisons. After controlling for abundance, only sapling growth potential, growth in shade and seedling mortality retain negative partial correlations, and part of that ordering may reflect the log-hazard scale itself, on which one diameter increment added to every sapling raises slow-growing species' promotion rates by the larger proportion.

covariate, mean over the interval's calendar years	ρ	P	ρ, windows one year later	P
Oceanic Niño Index, mean	−0.32	0.498	−0.50	0.267
fraction of months in El Niño conditions	+0.13	0.800	−0.36	0.429
fraction of months in La Niña conditions	+0.13	0.798	+0.29	0.556
dry-season rainfall	−0.57	0.200	−0.61	0.167
annual rainfall	−0.57	0.200	−0.18	0.713
solar radiation, January to April	−0.86	0.024	−0.61	0.167
solar radiation, annual	−0.79	0.048	−0.50	0.267
air temperature	−0.21	0.662	−0.18	0.713
dry-season length	+0.39	0.396	+0.29	0.556

Table S9: Rank correlations of the common sapling-growth variation with published climate series over the seven census intervals (descriptive only). The year effect here is the pooled sapling path (+0.35, +0.22, −0.26, −0.20, −0.10, −0.06, +0.05 for 1982–85 to 2010–15). ρ, Spearman rank correlation over the seven intervals with each covariate averaged over the interval's nominal calendar years (1982–84, 1985–89, …, 2010–14), and over windows set one year later because each census took about a year; P, two-sided, over all 5,040 orderings of seven values. The Oceanic Niño Index and its El Niño and La Niña classifications are the NOAA Climate Prediction Center series (14, 17, 20, 14, 13, 13, 6 El Niño months in the seven windows); rainfall, radiation, temperature and season dates are the Barro Colorado clearing-station summaries of Paton (2024), whose notes report a change of radiation sensor in March 1995 and of tower height in 2001. The largest $|\rho|$ of the nine, 0.86 on January-to-April radiation, has family-wise $P = 0.175$ (one-year-later windows: 0.61, $P = 0.70$); the path's correlation with the interval midpoint is −0.29 ($P = 0.56$). Seven serially correlated intervals can neither establish nor exclude a driver.

S4 Alternatives set aside

Purpose: three structural alternatives that we examined and did not adopt (a seedling class that is capped or omitted altogether, a capacity on basal area instead of stem number, and immigration weightings intermediate between the regional pool and exact compensation) and what each changes in the stationary community.

The five-year census enumerates only stems of at least 10 mm dbh and leaves open how the seedling class (class 1) should be limited. We compared three designs that differ only in that choice, each run to a stationary state at a common arrival rate and immigration weighting (36 replicate simulations from each of two families of initial conditions, Section S2). For each design Table S10 gives the stationary flow of stems from the seedling class into class 2 (14,900 to 35,300 per interval in the censuses; Table S1) and the stationary large-tree total. The first design imposes one capacity on all eight classes; because seedlings are below the census threshold, we computed the seedling component of that capacity from the fecundity rows of the projection matrices applied to the first-census counts in classes 2 through 8 (5,942,226 seedlings). Over a range of seedling capacities from one quarter to four times that value, with the demography unchanged, the stationary large-tree total scales almost in proportion to the capacity (log-log slope 0.986) while the rank correlation of predicted with observed abundance changes by only 0.04. Under this design the assumed seedling capacity therefore sets the large-tree total and has little effect on the predicted composition. At the computed value the flow into class 2 (38,786 stems per interval) exceeds the observed range. Of the capacities we tried, only half the computed value brings the flow into that range, and the large-tree total then falls to 0.54 times the observed value.

The second design, that of the main text, caps only the censused stand (classes 2 through 8) and lets births enter class 1 freely, so that the recruitment lottery for openings in the stand is the only density-regulated transition. With no cap of its own, the seedling class is still self-limiting: seed production by the standing forest balances the mortality of the seedlings that remain in class 1 while the stand is full. The flow into class 2 (34,117 stems per interval) falls inside the observed range, and the large-tree total is set by the demography and the immigrant stream rather than assumed (Results 3.2 of the main text, under both immigration schemes). The third design omits the seedling class, so that recruits enter class 2 directly. Most of them die before they grow, the large classes never fill, and the stationary large-tree total collapses to 27.9 stems, below the first design's by a factor of about 820. In contrast, the same omission changes a forecast two censuses ahead by at most 0.12% (Section S5).

In the main text the capacity is a number of censused stems, whatever their size; because larger stems use more space and resources, we also ran the model under the regional pool with the capacity placed instead on the basal area of the censused stand. Each class k is then weighted by the square of its midpoint diameter, which is $10^{0.6(k-2)}$ times the class-2 value and reaches 3,981 in the top class (closed at the next log-spaced edge, 1,259 mm). The basal-area capacity is the sum of the observed class totals with these weights, and in each interval the same recruitment lottery fills, in class-2 units, whatever the weighted stems remaining in the stand leave open. With the capacity on stem number, the stationary stand under the regional pool has 1.90 times the observed basal area, since it holds the observed number of censused stems with too many of them in the large-tree classes. Capped instead at the observed basal area, the stand halves (0.51 of the observed censused stems) and the large-tree total falls from 2.59 to 1.33 times the observed value. The composition barely changes (Table S11): the rank correlation differs by less than 0.01, and *Xylopia macrantha* remains the predicted dominant with 54% of the large-tree stand.

The regional pool and the compensated scheme are the two ends of a one-parameter family of

capacity design	seedling pool (stems)	flow into class 2 (stems per interval)	large trees (stems)	ratio to observed
(a) computed seedling capacity	5,915,654	38,786	22,902	1.06
(b) uncapped seedling class	5,372,837	34,117	20,245	0.93
(c) no seedling class	0	(none)	27.9	0.0013

Table S10: [**PLACEHOLDER: this table is to be regenerated before submission with arrivals weighted by regional share at 73,729 immigrant seeds per interval, the regional pool of the main text; the values shown come from simulations that weighted arrivals by each species' published demographic variance at 147,260 immigrant seedlings per interval, a weighting we do not otherwise use, and the text draws only on the contrasts between the three designs, which do not depend on it.**] Stationary states of the three capacity designs, which differ only in what limits the seedling class; all three use the published projection matrices of the 84 species, the same immigrant stream and 36 replicate simulations from each of two families of initial conditions. Seedling pool: the stationary number of stems below 10 mm. Flow into class 2: stems promoted from class 1 into the smallest censused class per five-year interval; the census data give 14,900 to 35,300. Large trees: stationary stems of at least 79.4 mm dbh (classes 5 through 8), and their ratio to the observed 21,693. Design (a) imposes a computed seedling capacity of 5,942,226 stems, of which the pool shown is the stationary occupancy; design (b), the model of the main text, leaves the seedling class uncapped, and under the regional pool its large-tree total is that of Results 3.2 of the main text; design (c) has no seedling class and recruits enter class 2 directly. In these simulations only (b) gives both a flow inside the observed range and a large-tree total near the observed one.

capacity on	censused stand (× obs.)	basal area (× obs.)	large trees (× obs.)	ρ	R^2	dominant	*Xylopia* share (%)	below one (species)
stem number	1.00	1.90	2.59	0.667	−0.741	*X. macrantha*	54	17
basal area	0.51	0.99	1.33	0.665	−1.053	*X. macrantha*	54	23

Table S11: The stationary community under the regional pool with the capacity on stem number, as in the main text, and with the capacity on basal area; 73,729 immigrant seeds per interval weighted by regional share, 36 replicate simulations from each of two families of initial conditions, means over the final analysis windows of Section S2. The simulations give the nine species absent from the regional plots zero weight, and the stem-number row therefore differs slightly from the closed-form values of the main text. Censused stand: stationary stems in classes 2 through 8 relative to the observed 200,721. Basal area: the stand's basal area relative to the observed stand's. Large trees: stationary stems in classes 5 through 8 relative to the observed 21,693. ρ and R^2: Spearman rank correlation, and identity-line R^2 on log abundance, of stationary against census-mean large-tree abundance over the 84 species. Under both capacities the observed dominant, *Faramea occidentalis*, is not the most abundant predicted species. Below one: species whose stationary large-tree abundance is below one stem (5 observed).

weightings, $w_s \propto p_s / G_s^{\zeta}$, in which arrivals are weighted by regional share p_s and against a power of the local gain $G_s = (1-d)\,\Omega_s/(d\,\mu_s)$, the large trees standing per arriving seed. At $\zeta = 0$ the weighting is the regional pool, and at $\zeta = 1$ it is the compensated scheme, a competition-colonization trade-off in its simplest form (Levins and Culver, 1971; Hastings, 1980; Tilman, 1994). Under that scheme the gain cancels from the expected large-tree abundance, and the composition is the regional composition p_s itself. We computed the stationary community along this family from the closed form of Section S2 at 73,729 immigrant seeds per interval (Table S12).

As the exponent ζ increases from 0 (the regional pool) to 1 (the compensated scheme), the large-tree total falls from 2.59 to 1.25 times the observed value. Over the same range the number of species below one stem falls from 16 to 5 (the observed number), *X. macrantha*'s share shrinks, and the identity-line R^2 on log abundance improves. The rank correlation changes least, from 0.67 to at most 0.731 and hardly at all beyond $\zeta = \frac{1}{2}$, and *Faramea occidentalis*, the observed dominant, rises only to rank 10. Past exact compensation, at $\zeta = 1\frac{1}{4}$, the large-tree total falls below the observed one with no further improvement in R^2. Because the intermediate weightings interpolate between the two ends, we present only the ends in the main text, as distinct hypotheses about how immigrants establish in this forest (Fig. S2). Equalizing the gains by adjusting each species' fecundity, instead of weighting its arrivals against the gain as in the compensated scheme, gives the same expected large-tree composition (Results 3.2 of the main text), but the fecundities required are not feasible. Even with no local reproduction, each immigrant seed of *X. macrantha* yields 6.0 times as many large trees as a seed of the species with the median gain, and the nearest feasible adjustment gives a rank correlation of 0.723 and $R^2 = +0.463$.

	rank correlation ρ			R^2 of log	large trees	*Xylopia*	*Faramea*	below
ζ	full	rare	common	abundance	(× obs.)	share (%)	rank	one stem
0 (regional pool)	0.666	+0.66	+0.13	−0.693	2.59	54.2	23	16
$\frac{1}{4}$	0.698	+0.69	+0.22	−0.243	2.06	33.7	21	13
$\frac{1}{2}$	0.723	+0.72	+0.29	+0.100	1.73	17.1	16	10
$\frac{3}{4}$	0.731	+0.74	+0.32	+0.346	1.49	7.3	14	9
1 (compensated scheme)	0.727	+0.74	+0.40	+0.495	1.25	2.6	10	5
$1\frac{1}{4}$	0.718	+0.74	+0.44	+0.485	0.87	0.8	10	5

Table S12: Arrivals weighted by regional share and against the local gain, $w_s \propto p_s / G_s^{\zeta}$, from the regional pool ($\zeta = 0$) to the compensated scheme ($\zeta = 1$, whose expected large-tree composition is the regional composition) and beyond, computed from the closed form of Section S2 at 73,729 immigrant seeds per interval, with the nine species absent from the regional plots entered at the smallest regional share. Rank correlation: Spearman correlation of stationary with census-mean large-tree abundance (stems of at least 79.4 mm dbh) over all 84 species and over the rarer and the commoner half, split at the median observed abundance. R^2: identity-line R^2 on log abundance. Large trees: the stationary large-tree total relative to the observed 21,693. The last three columns give *Xylopia macrantha*'s share of the large-tree stand, the predicted rank of *Faramea occidentalis* (observed first) and the number of species below one stem (5 observed).

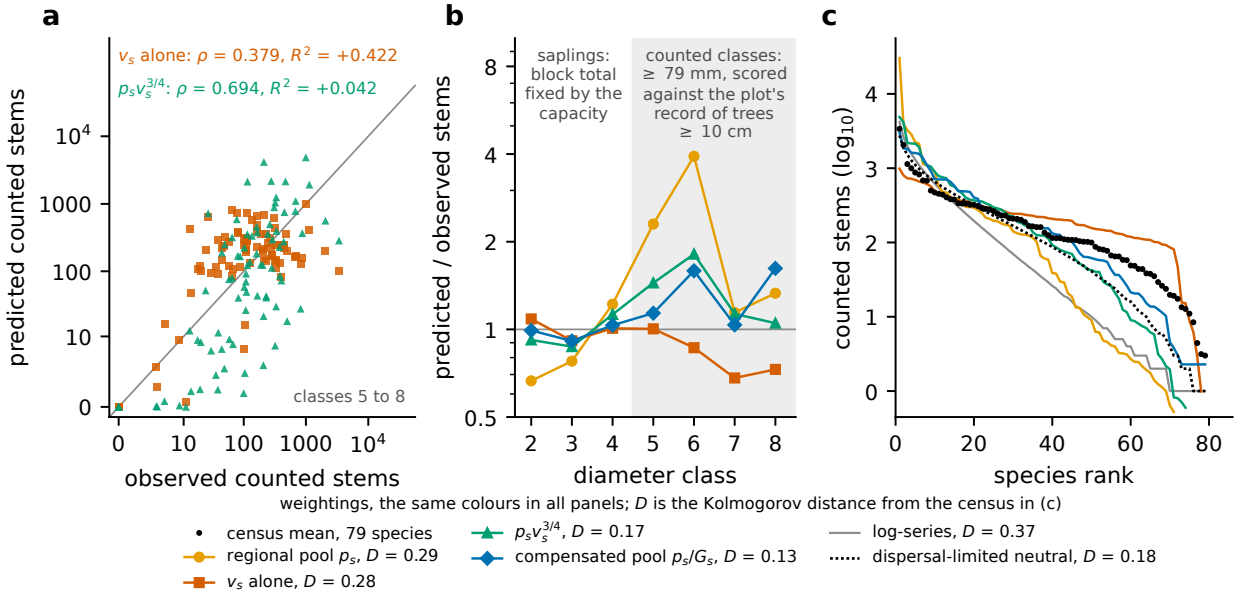

Figure S2: [**PLACEHOLDER: this figure is to be redrawn before submission from the same closed-form results with the two demographic-variance weightings removed; the caption describes the redrawn figure, and the file shown still draws those weightings and a first panel that the redraw drops.**] The distribution of stems across size classes and across species under the two immigration schemes of the main text. (a) Stationary stems in each censused diameter class, classes 2 through 8, relative to the observed number, under the regional pool (circles) and the compensated scheme (diamonds). The capacity holds the censused-stand total at the observed value under both schemes, so an excess in the large-tree classes (shaded; stems of at least 79.4 mm dbh) is offset by a deficit in the sapling classes. (b) The abundance distribution of the large-tree stand: ranked $\log_{10}$ abundance of the 79 species that reach the large-tree classes in the census (points), under the regional pool and the compensated scheme (Kolmogorov distances from the census on log abundance 0.290 and 0.127), with a log-series (gray; 0.37) and Hubbell's dispersal-limited neutral model fitted to the census abundance distribution of the same species (dotted; 0.18) for comparison. Closed form throughout, with the nine species absent from the regional plots entered at the smallest regional share; intermediate weightings are in Table S12.

S5 The decadal test in detail

Purpose: the decade-ahead test of the main text (its Section 3.1 and Fig. 3) in full: the summary table of predictors the main text cites, the comparators, every division of the censuses into an estimation window and tested censuses, changes of direction, per-species net growth rates, the seedling class and the community total.

Throughout, the quantity we predict is a species' change in large trees (stems ≥ 79.4 mm dbh, classes 5 to 8 of the published matrices) one and two censuses after a starting census, the $\Delta_5 N_s$ and $\Delta_{10} N_s$ of the main text (Material and methods 2.6). We obtain it by applying the species' projection matrix, the five-year operator A_s (Jops et al., 2025), once or twice to the observed stage vector at the start, summing over the large-tree classes and subtracting the starting number; a species-start pair is one species at one starting census. The horizons are nominal, because the first interval (1982 to 1985) was three years. To a good approximation only the survival and growth entries of the censused classes matter at this horizon, because under every published matrix a seedling contributes at most 0.0003 stems to the large-tree classes two censuses later and a stem in the smallest censused class at most 0.013 per census. The capacity, the recruitment lottery and immigration therefore do not enter the projection, and we hold the year effect at its mean. Passing the same starting vectors through the full model's mean dynamics instead, with the seedling stage, each species' fecundity row, immigration and the capacity acting, leaves the one-census predictions identical and moves the two-census predictions by 0.44 stems per species-start pair in sample and 0.48 out of sample (1.6% and 1.5% of the mean predicted change), the two sets of predictions having rank correlations of 0.998 and 0.999; the forecast therefore tests the survival and growth structure of the censused stems and cannot distinguish our model from any other that shares those rates, the mechanisms of long-run abundance being tested in Section S7.

Table S13 summarizes the test as the main text reports it, the projection matrices against the two comparators without demography, in sample and out of sample; the paragraphs and tables that follow give each part in full.

Comparators inside the census series. Within the census series of 1982–2015, from which Jops et al. (2025) estimated the operators, we apply each species' operator at each of the first five censuses (420 species-start pairs). On the same pairs we compare it with predictors that use none of the species' own demographic rates (Table S14), among them the species' trend over the intervals outside the tested window and its momentum (the change over the preceding interval, continued). The operators' rank correlation between predicted and observed change, 0.81 one census ahead and 0.86 two censuses ahead, exceeds that of the other-interval trend by +0.16 (90% CI over species +0.10 to +0.24) and +0.21 (+0.14 to +0.32). The predicted magnitudes agree with the observed about as well as the ranks do. When we express both as net growth rates, weight each pair by its large trees at the start and remove the community-wide mean shift, the predicted rates account for 0.748 and 0.848 of the variance of the observed five- and ten-year rates about the identity line.

Between species (each species' mean predicted against its mean observed change over the five starts) the rank correlation is 0.96 and 0.98: being pooled transition frequencies, the operators reproduce each species' average trend over the censuses they were estimated from. Within species (each start's deviation from the species' own mean) it is 0.28 and 0.29, and computed for each species separately over its five starts it is positive in 62 of the 79 species whose numbers vary. This within-species correlation measures how well a predictor distinguishes the starts at which a species runs ahead of its own average trend from those at which it falls behind. The operators make that distinction through the stage structure at each start, and an extrapolated trend, which predicts the same change at every start, cannot.

Table S13: Decade-ahead change of each species' large trees, predicted from stage-structured rates and from two comparators with no demography, in sample and out of sample. A species-start pair is one species at one starting census. Sign, percentage of pairs with nonzero observed change predicted in the right direction; ρ, Spearman rank correlation over pairs (90% CI over species); within, rank correlation of each start's deviation from its species' mean (undefined for a trend); MAE, mean absolute error in stems. Trend, each species' mean change per interval from the intervals outside the predicted window (in sample) or from 1982–1995 (out of sample); momentum, its change over the previous interval, defined from the second census on; both doubled at two censuses. A reversal is an observed change one census ahead opposite in sign to the species' net change from 1982 to the start, counted in sample at every start from 1985 to 2010 (504 pairs), where the matrices' rates include the predicted intervals; a trend fitted up to the start anticipates none, and the same count with matrices estimated without the predicted interval is in Table S18. Predicting no change gives MAE 16.1 and 30.4 stems; further comparators are in Table S14 and other early–late splits of the censuses in Table S16.

		one census ahead				two censuses ahead				
predictor	pairs	sign (%)	ρ (90% CI)	within	MAE	sign (%)	ρ (90% CI)	within	MAE	reversals anticipated
In sample: rates from 1982–2015; starts 1982–2000										
projection matrices	420	84	0.81 (0.75–0.87)	0.28	8.0	85	0.86 (0.80–0.91)	0.29	12.1	78 of 128 (61%, 90% CI 52–69%); trend up to the start, none
trend, other intervals	420	74	0.66	–	10.0	71	0.65	–	18.7	–
momentum	336	72	0.67	–	11.2	72	0.64	–	21.0	–
Out of sample: rates from 1982–1995 only; starts 1995, 2000 and 2005										
projection matrices	252	74	0.71 (0.60–0.80)	0.47	11.4	71	0.68 (0.55–0.79)	0.46	22.1	–
trend, 1982–1995	252	69	0.54	–	11.8	66	0.51	–	23.2	–
momentum	252	74	0.64	−0.12	8.3	71	0.62	−0.23	16.4	–

		one census ahead			two censuses ahead		
predictor	pairs	sign (%)	ρ	MAE	sign (%)	ρ	MAE
projection matrices (published)	420	84	0.81	8.0	85	0.86	12.1
species trend, other intervals	420	74	0.66	10.0	71	0.65	18.7
momentum	336	72	0.67	11.2	72	0.64	21.0
mean operator	420	49	0.02	60.1	46	0.01	100.4

Table S14: Each species' published projection matrix (the five-year operator A_s) inside the census series of 1982–2015 compared with predictors that use none of the species' own demographic rates, on the same species-start pairs (84 species at each of the first five censuses) and in the same way, one and two censuses ahead. Sign: the percentage of pairs with a nonzero observed change in large trees (stems ≥ 79.4 mm dbh) for which the predicted direction is correct; ρ: the Spearman rank correlation over pairs; MAE: the mean absolute error in stems. Projection matrices (published): the decadal change $\Delta_5 N_s$, $\Delta_{10} N_s$ defined in Material and methods 2.6 of the main text. Species trend: each species' mean observed change per interval over the intervals outside the tested window, doubled at two censuses. Momentum: the species' observed change over the interval before the start, doubled at two censuses, defined from the second census onward (on those 336 pairs the operators give 0.84 and 0.88). Mean operator: the average of the 84 published operators applied to each species' own observed stage vector. Predicting no change has a mean absolute error of 16.1 and 30.4 stems. The operators were estimated from stems observed in the tested intervals; the other-interval trend and momentum were not.

The forecast out of sample. The forecast proper puts the model in the position of most forest plots, which have only a few censuses, and asks what could have been predicted from them at the time. Matrices tabulated from the whole series differ from such early ones in resolving each species' life history more finely, and Table S16 below shows how the skill changes with the censuses used, including the case in which only the predicted interval is withheld; the more of the life history the censuses have measured, the better the forecast. For the forecast proper we estimated each species' survival and growth frequencies among classes 2 to 8 afresh from the public tree tables of the first four censuses alone (1982–1995), following the procedure and stage boundaries of Jops et al. (2025). We applied them at the censuses of 1995, 2000 and 2005 and compared the projections with the censuses through 2015 (252 species-start pairs at each horizon; Table S15 and Fig. S3). Estimated from the whole series in the same way, the frequencies reproduce the published fits' rank correlations and sign agreement (in-sample rows). Over the single decade 1995–2005 with all 84 species (third row) the rank correlation is similar, but the forecast is biased upward: it lies above the observed change for 65 species, and the predicted changes sum to +1,167 stems, compared with −578 observed. On the two late targets of the in-sample rows the rank correlation of the whole-series rates exceeds that of the 1982–2005 rates by +0.18 and +0.22, the contribution of the tested intervals' own transitions.

The 1982–1995 rates also preserve the within-species timing out of sample: the rank correlation of each start's deviation from the species' own mean is 0.47 (90% CI 0.33 to 0.60) at one census and 0.46 at two. A continued trend has no such correlation, and that of momentum is negative, −0.12 (−0.23 to +0.01) and −0.23, because a species' larger changes tend to be followed by smaller ones. Momentum's overall rank correlation of 0.64 therefore comes entirely from ordering the species against one another (between-species correlation 0.88), although its absolute error is smaller than that of the species' own operators at both horizons. We conclude that the within-species timing comes from each species' stage structure at the start and the between-species ordering from its

own rates, and that only the species' own operators combine the two. Dropping the three-year first interval from the estimation changes the rank correlations of the forecast by −0.03 and −0.02. Partial pooling of the early operators toward the community's changes little at this plot: shrinking each species' 1982–1995 transition frequencies toward the pooled frequencies of all 84 species, with weight $n/(n+\hat{\tau})$ on a species' own n stems in a class and $\hat{\tau}$ between 9 and 28 stems by the method of moments, gives rank correlations of 0.72 and 0.69, a change of +0.01 (90% CI -0.01 to +0.02) at one census, because every species followed has at least 774 stem-intervals in those censuses; heavier pooling toward the community operator lowers the skill. At sparser plots such pooling would be the natural default.

Every contiguous division of the censuses. In Table S16 we repeat the test for every contiguous way of dividing the censuses of 1982–2015 into an estimation window and tested censuses, to check that the result does not depend on the choice of 1982–1995 as the estimation window. Forward rows use the earliest intervals and backward rows the latest, each with its own window's trend as comparator. In the 18 forward and backward cases with at least two intervals in the window the rank correlation over species-start pairs lies between 0.60 and 0.73 and exceeds that of the window's own trend by +0.05 to +0.22. The backward cases do as well as the forward ones. A single interval, or the longest forward window with only the start of 2010 left to test, gives the lowest values. Leaving out only the window a pair is tested on (the last row) gives 0.75 and 0.77 over the 420 pairs, with the between-species order intact (0.95 and 0.97). The within-species part is then negative (−0.05 and −0.20), as expected when the transitions removed are precisely those of the tested window. In every division with at least two starts and two intervals in the window the operators' within-species correlation exceeds that of momentum, and the 90% interval over species on the difference excludes zero.

Changes of direction. Aggregate rank correlations do not show how often a predictor anticipates a change of direction. A continued trend cannot reverse, but a projection of the stage structure at the start under the species' fixed rates can predict a change opposite in sign to that trend. Of the 252 species-start pairs tested out of sample (Table S15), 66 are reversals at one census, the large trees changing after 1995 in the direction opposite to the species' own change over 1982–1995, and 73 at two censuses. At one census another 148 are continuations, and in the remaining 38 pairs either change is zero (Table S17). From the stage structure at the start alone, the rates measured up to 1995 anticipate 19 of the 66 reversals, 29% (90% CI over species 16 to 42%), and 21 of the 73 at two censuses. They give the correct direction for 140 of the 148 continuations.

Momentum repeats the sign of the change over the interval before the start, a change that the projection from the starting stage vector does not use. It therefore anticipates a reversal only when that reversal has already begun in the preceding interval, and it misses a continuation wherever the change over that one interval was opposite in sign to the species' earlier change. Across the 18 divisions of Table S16 the operators anticipate between 21 and 35% of the reversals defined against their own estimation windows. Table S18 gives the comparison of the main text (Results 3.1), made in sample at every starting census from 1985 to 2010 with a reversal counted as a change opposite in sign to the species' net change from 1982 to the start, together with its out-of-sample check, the same rates re-estimated with the tested (focal) interval left out. The published operators, estimated from all seven intervals including those tested, anticipate 78 of the 128 reversals one census ahead, 61% (90% CI over species 52 to 69%), and 70 of 110 two ahead, and give the correct direction for 288 of the 306 continuations; the trend up to each start anticipates no reversal and every continuation. Re-estimated without the tested interval, out of sample everywhere while using

rates estimated from	tested on	pairs	sign (%)	ρ	trend ρ	within	MAE
1982–1995	1995, 2000, 2005 → one ahead	252	74	0.71	0.54	0.47	11.4
1982–1995	the same → two on	252	71	0.68	0.51	0.46	22.1
1982–1995	1995 → 2005	84	77	0.75	0.59		22.1
1982–1995	2005 → 2010	84	69	0.63	0.45		11.4
1982–2005	2005 → 2010	84	69	0.62	0.50		8.6
1982–2005	2005 → 2015	84	66	0.61	0.51		16.9
1982–2005	2010 → 2015	84	70	0.56	0.51		10.1
1982–2010	2010 → 2015	84	74	0.63	0.59		9.3
1982–2015 (in sample)	2005 → 2010	84	80	0.80	0.68		6.4
1982–2015 (in sample)	2010 → 2015	84	82	0.78	0.71		8.0
published fits (in sample)	2005 → 2010	84	80	0.80	0.68		6.4
momentum	1995, 2000, 2005 → one ahead	252	74	0.64		−0.12	8.3
community operator, 1982–1995	the same	252	53	0.26		0.40	21.3

Table S15: Rates estimated from part of the census series and tested on the rest, beside the same rates estimated from all of it. Each species' survival and growth frequencies among classes 2 to 8 are estimated from the public tree tables of the censuses named in the first column, with the stage boundaries of the published fits, and applied to the species' observed stage vector at the starting census. The predicted change in its large trees (stems ≥ 79.4 mm dbh) is compared with the census named in the second column, one or two censuses after the start, one species-start pair per species and start. Published fits: the operators of Jops et al. (2025) themselves. Sign: the percentage of pairs with a nonzero observed change for which the predicted direction is correct. ρ: the Spearman rank correlation over pairs. Trend ρ: the same for each species' mean change per interval over the estimation window continued forward, the only comparator without stage-structured rates that uses the same censuses as the estimated operators and no others. Within: the rank correlation of each start's deviation from the species' own mean over its starts (defined only for several starts; a continued trend has none, since it predicts the same change at every start). MAE: the mean absolute error in stems. Momentum: the change over the interval before the start. Community operator: a single operator estimated from the pooled stems of all species, applied to each species' own stage vector. Above the first rule no tested interval enters the estimation window; in the in-sample rows it does (see text). For most species the change over 2010 to 2015 continued that over 2005 to 2010 (rank correlation 0.756 between the two), a degree of persistence that rates averaged over earlier intervals could not be expected to anticipate, and the 2010 to 2015 target accordingly has the lowest ρ for every estimation window that ends before 2010. The first two rows and the momentum row are the out-of-sample rows of Table 2 of the main text.

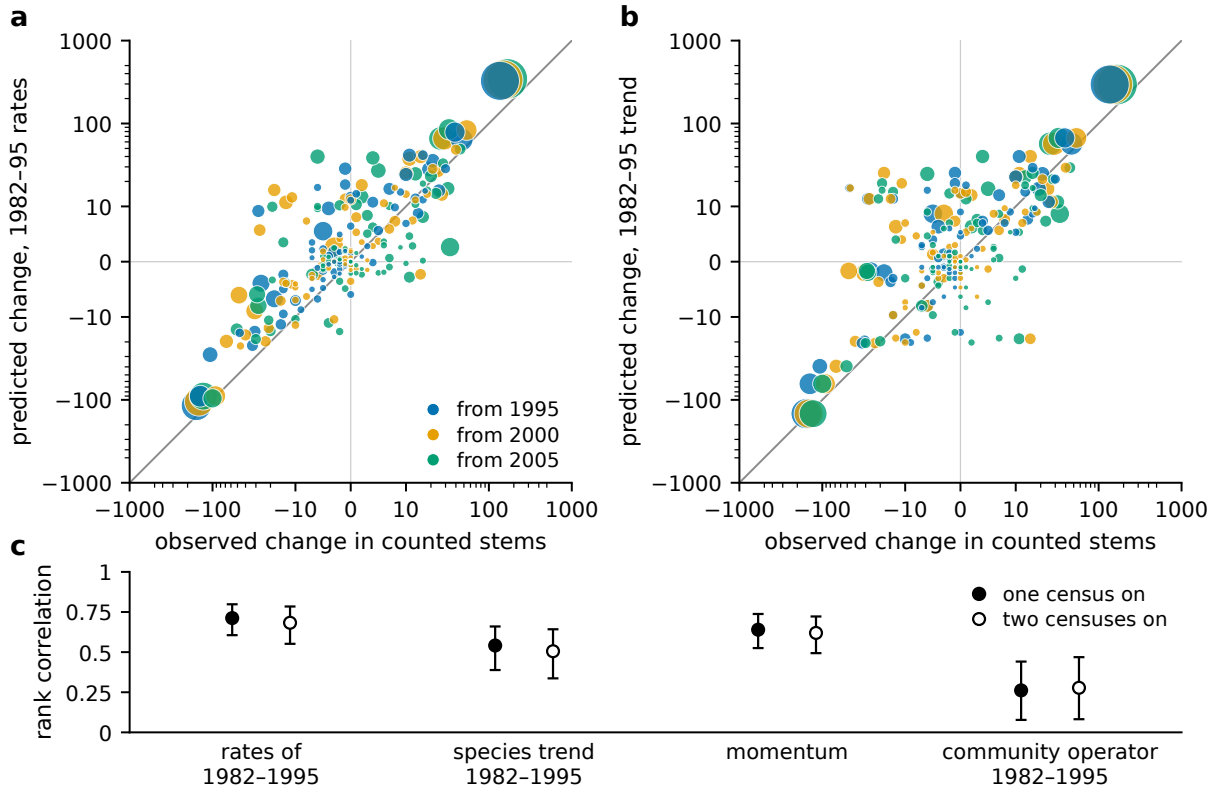

Figure S3: The out-of-sample test pair by pair, one census ahead. (a) Predicted versus observed change in a species' large trees (stems ≥ 79.4 mm dbh, classes 5 to 8) one census after the start; the prediction is the species' survival and growth rates estimated from the 1982–1995 censuses applied to its observed stage vector at the censuses of 1995, 2000 and 2005. One point per species and starting census (252 points, 84 species), colored by starting census, with area proportional to the species' large trees at the start; the line is the identity and the axes are linear within ten stems and logarithmic beyond. (b) The same pairs on the same axes when each species is given its mean change per interval over 1982–1995. (c) Rank correlation over the 252 pairs for the four predictors of Table S15, one census ahead (filled) and two censuses ahead (open), with 90% intervals over species. One census ahead the rank correlation is 0.71 (0.60 to 0.80) for the early rates and 0.54 for the trend, a margin of +0.17 (+0.09 to +0.27). In absolute numbers the forecasts are biased high: 176 of the 252 points in (a) lie above the identity line, and the predicted one-census changes sum over species to +1,729 stems, compared with -517 observed. The two-census counterpart of (a) is Fig. 3b of the main text.

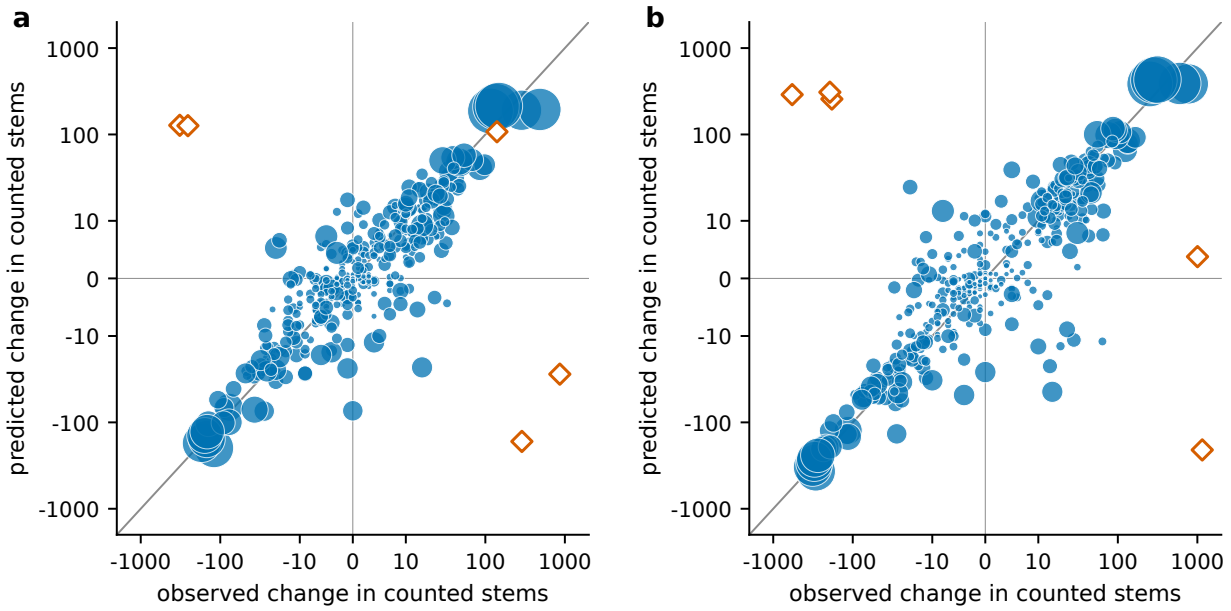

Figure S4: Predicted versus observed decadal change inside the census series (the within-series consistency test of the main text), pair by pair. Each point is the change in one species' large trees (stems ≥ 79.4 mm dbh) from one of the first five censuses, 1982 to 2000 (420 points, 84 species), one census ahead in (a) and two censuses ahead in (b). The area of a point is proportional to the species' large trees at the start, the line is the identity, and the axes are linear within ten stems and logarithmic beyond. The prediction is the species' published operator (Jops et al., 2025), estimated from all the censuses of 1982–2015, applied once or twice to its observed stage vector; nothing else in the model enters. The rank correlation between predicted and observed change is 0.81 and 0.86 (90% intervals over species 0.75 to 0.87 and 0.80 to 0.91). The predicted sign is correct in 312 of the 371 pairs with a nonzero observed change at one census and 326 of 382 at two, and the mean absolute error is 8.0 and 12.1 stems per pair (median over species 4.7 at one census). Restricted to species with at least 10 large trees at the start (373 pairs of 76 species) the rank correlations are 0.83 and 0.87. Open diamonds are the five community totals, the sums over species from each starting census (Table S20), which the summed projections do not reproduce. Panel (b) is Fig. 3a of the main text; comparators are in Table S14.

Table S16: The decadal test under every contiguous division of the censuses. Rates estimated from: the censuses whose intervals supply the pooled transition frequencies among classes 2 to 8 (as in Table S15), and the number of intervals. For one census ahead and for two: the starting censuses tested (every census outside the estimation window whose target census exists), the species-start pairs (84 per start), the rank correlation ρ between predicted and observed change in large trees over those pairs, the same for the window's own trend (each species' mean change per interval over the estimation window, times the horizon), and the difference. Forward rows test later censuses with earlier rates, the first one to six intervals top to bottom; backward rows test earlier censuses with later rates, the last one to six intervals, the prediction running forward from the start. The first row uses every interval and is the within-series consistency test of Table S14; a trend over those same intervals would contain the tested change itself and is not entered, its comparators being those of Table S14. The last row leaves out of the estimation only the interval a pair is tested on (the two intervals, at two censuses ahead), and its trend likewise leaves that window out. †One interval in the estimation window. A dash under two censuses ahead marks a division with no census two steps beyond its last start.

	rates estimated from		one census ahead					two censuses ahead				
	censuses	int.	starts	pairs	ρ	trend	diff.	starts	pairs	ρ	trend	diff.
all intervals	1982–2015	seven	1982–2000	420	0.81	–	–	1982–2000	420	0.86	–	–
forward†	1982–1985	1	1985–2010	504	0.55	0.45	+0.10	1985–2005	420	0.60	0.49	+0.11
	1982–1990	2	1990–2010	420	0.66	0.54	+0.12	1990–2005	336	0.68	0.54	+0.14
	1982–1995	3	1995–2010	336	0.66	0.51	+0.15	1995–2005	252	0.68	0.51	+0.18
	1982–2000	4	2000–2010	252	0.64	0.52	+0.12	2000, 2005	168	0.66	0.52	+0.14
	1982–2005	5	2005, 2010	168	0.60	0.51	+0.09	2005	84	0.61	0.51	+0.10
	1982–2010	6	2010	84	0.63	0.59	+0.05	–	–	–	–	–
backward†	2010–2015	1	1982–2005	504	0.61	0.51	+0.10	1982–2000	420	0.65	0.53	+0.12
	2005–2015	2	1982–2000	420	0.61	0.46	+0.16	1982–1995	336	0.65	0.49	+0.16
	2000–2015	3	1982–1995	336	0.66	0.49	+0.16	1982–1990	252	0.68	0.53	+0.15
	1995–2015	4	1982–1990	252	0.68	0.47	+0.22	1982, 1985	168	0.68	0.50	+0.19
	1990–2015	5	1982, 1985	168	0.71	0.53	+0.18	1982	84	0.73	0.59	+0.14
	1985–2015	6	1982	84	0.67	0.53	+0.14	–	–	–	–	–
all but the tested window	1982–2015	six (five)	1982–2000	420	0.75	0.66	+0.10	1982–2000	420	0.77	0.65	+0.12

predictor	reversals anticipated		continuations anticipated	
	one census ahead (of 66)	two censuses ahead (of 73)	one census ahead (of 148)	two censuses ahead (of 148)
rates from 1982–1995	19 (29%)	21 (29%)	140 (95%)	139 (94%)
species trend, 1982–1995	0	0	148	148
momentum	34 (52%)	32 (44%)	124	124

Table S17: Changes of direction, out of sample. The 252 species-start pairs of Table S15 (84 species starting at the censuses of 1995, 2000 and 2005) split by whether the observed change in large trees over the tested window has the sign opposite to the species' own change over 1982 to 1995 (a reversal) or the same sign (a continuation); pairs in which either change is zero (38 at one census, 31 at two) are in neither column. An entry is the number of those pairs whose direction the predictor anticipates, the predictors being those of Table S15. A continued trend cannot reverse and therefore anticipates no reversal and every continuation. The 90% interval over species on the rates' reversal fraction is 16 to 42% at one census and 17 to 41% at two. Requiring a change of at least 2 stems on both sides leaves 41 and 63 reversals, of which the rates anticipate 11 and 18. Over all pairs with a nonzero observed change the rates give the correct direction in 74 and 71% of cases.

nearly the whole series, the same rates anticipate 66 of the 128 reversals one census ahead, 52% (90% CI over species 43 to 60%), and 53 of 110 two ahead, and give the correct direction for 278 of the continuations: 9 points of the reversal fraction (4 to 15 over species) come from rates that include the tested interval, and about half the reversals are anticipated by rates that do not, which suggests that reversals arise partly from the stage structure and rates at the start and partly from change in the rates between intervals. The reversals that a fixed-rate projection from the starting stage structure does not anticipate presumably reflect changes in the rates themselves within the tested window.

Per-species net growth rates. A trend supplies each species' net growth rate as well as an ordering of species, and attaching that rate to the early operators does not improve the forecast, because they already contain it. Being pooled transition frequencies, they recover approximately the observed survivors of the 1982–1995 intervals class by class when applied to the stage vectors of those censuses, and so reproduce, very nearly, each species' observed net change in large trees over that window. The ratio γ_s of a species' observed five-year growth factor in large trees over those intervals to the factor its operator implies has median 1.00 over the 84 species and range 0.96 to 1.10. The normalization of every published matrix to growth rate one does not remove this observed net growth from a projection of one or two censuses, because the normalization acts through the fecundity row.

In Table S19 we nevertheless compare two ways of attaching a per-species net growth rate to the early operators. Multiplying each early operator by γ_s, compounded as γ_s^h at h censuses, leaves the out-of-sample rank correlation almost unchanged (paired difference −0.005, 90% CI −0.014 to +0.003). The operators scaled in this way have a slightly larger absolute error, lower sign agreement and fewer reversals anticipated, as expected of a predictor that gives extra weight to each species' earlier direction of change. Because the survival-and-growth part of each matrix alone has a leading eigenvalue of median 0.90, rescaling each operator until its leading eigenvalue equals the observed early growth factor instead inflates every projection by a factor of about 1.11. The rescaled operators then predict an increase for nearly every pair (rank correlation 0.34 at

predictor	reversals anticipated		continuations anticipated	
	one census ahead (of 128)	two censuses ahead (of 110)	one census ahead (of 306)	two censuses ahead (of 261)
published fits (in sample)	78 (61%)	70 (64%)	288 (94%)	245 (94%)
rates without the tested interval (out of sample)	66 (52%)	53 (48%)	278 (91%)	235 (90%)
rates from 1982–1995	49 (38%)	42 (38%)		
species trend up to the start	0	0	306	261
momentum	41 (32%)	31 (28%)	270	234

starting census (one census ahead)	1985	1990	1995	2000	2005	2010
reversals	26	16	17	23	26	20
anticipated by the published fits (in sample)	17	12	10	16	13	10
by the rates without the tested interval	16	11	9	12	11	7
by the rates from 1982–1995	16	6	5	7	9	6

Table S18: Changes of direction from every starting census, with the published fits in sample. All 84 species at every starting census that has an earlier interval: 1985 to 2010 one census ahead (504 species-start pairs, of which 444 changed) and 1985 to 2005 two censuses ahead (420 pairs, 380 changed). A reversal is a pair whose observed change in large trees (stems ≥ 79.4 mm dbh) has the sign opposite to the species' net change from 1982 to the start, a continuation one with the same sign; the 70 and 49 pairs with no change on one side are in neither column. Published fits: the survival and growth entries of the matrices of Jops et al. (2025), estimated from all seven intervals of 1982 to 2015, so that every tested interval lies inside their estimation window. Rates without the tested interval: the same entries re-estimated from the six other intervals (five, two censuses ahead), out of sample everywhere. Rates from 1982 to 1995: the operators of Table S17 applied at every start; on the starts from 1995 on, where they are out of sample, they anticipate 27 of 86 reversals (31%), compared with 49 (57%) for the published fits and 39 (45%) for the rates without the tested interval; their continuation entries over all starts are omitted because two of the six starts lie inside their window. Lower table: the reversals from each starting census one census ahead and the number each set of rates anticipates. The trend, each species' mean change per interval from 1982 to the start continued, cannot reverse and so anticipates no reversal and every continuation; momentum repeats the sign of the last interval. The 90% interval over species on the published fits' reversal fraction is 52 to 69% one census ahead and 53 to 73% two ahead; on that of the rates without the tested interval, 43 to 60% and 36 to 59% (a difference of 9 points one census ahead, 4 to 15). Defining a reversal instead against the change over the last interval alone gives 106 and 90 reversals, of which the published fits anticipate 63 and 67%, the rates without the tested interval 57 and 51%, and the trend 25 and 20%; requiring at least 2 stems of change on both sides leaves 73 reversals one census ahead, of which the published fits anticipate 59% in sample.

Table S19: Per-species net growth rates attached to the early operators, out of sample. Pairs, statistics, the early operators, the species trend and momentum as in Table S15 (rates from 1982 to 1995; 84 species starting at the censuses of 1995, 2000 and 2005). Operators $\times \gamma_s$: each species' early operator multiplied by the ratio γ_s of its observed to implied five-year growth factor in large trees over 1982 to 1995 (median 1.00, range 0.96 to 1.10), applied h times. Operators, eigenvalue: each operator rescaled until its leading eigenvalue equals the observed early growth factor. Within: as in Table S15. Reversals (rev.) and continuations (cont.): as in Table S17 (of 66 and 148 at one census, 73 and 148 at two). Sign in %. Lower table: the 1995 stage vectors projected four intervals to 2015; total large trees of the 84 species predicted for 2015 (observed 21,859; 22,244 in 1995), mean absolute error per species, rank correlation and sign agreement of the twenty-year changes, and species whose predicted number is off by more than a factor of two.

	one census ahead						two censuses ahead			
predictor	sign	ρ	within	MAE	rev.	cont.	sign	ρ	MAE	rev.
early operators	74	0.712	0.47	11.4	19	140	71	0.683	22.1	21
operators $\times \gamma_s$	72	0.707	0.45	11.9	15	141	71	0.684	22.9	18
operators, eigenvalue	48	0.34		46.0	8	96	50	0.33	101.5	14
species trend	69	0.54	–	11.8	0	148	66	0.51	23.2	0
momentum	74	0.64	−0.12	8.3	34	124	71	0.62	16.4	32

four censuses ahead (1995 to 2015)	total 2015	MAE	ρ of change	sign	off twofold
early operators	24,556	41.7	0.71	68	3
operators $\times \gamma_s$	24,809	43.5	0.70	66	2
species trend, four intervals	23,967	43.7	0.52	68	10
no change	22,244	52.8	–	0	2

one census). When the 1995 stage vectors are projected four intervals forward, to 2015, the early operators predict 24,556 large trees for the modeled species and the four-interval trend 23,967, compared with 21,859 censused (Table S19, lower part). Both predictors overstate the censused total, which had fallen slightly over those two decades. A net rate measured over one decade need not persist, and in the stationary model of the main text we therefore hold every species at growth rate one and let immigration set absolute abundances.

The seedling layer and the community total. In the decadal projections we complete the seedling entry of each observed stage vector, which the census does not enumerate, by applying the species' published fecundity row to its censused counts. Deleting the seedling class instead (the variant of Section S4) makes almost no difference over a decade. Across the five starting censuses the two variants' predictions of the large-tree total differ by between 11.2 and 27.1 stems (mean 21.5), at most 0.12% of the observed total at the start, and the largest difference for any single species-start pair is 17 stems. The two-census rank correlation between predicted and observed species changes is 0.86 with the seedling class and 0.86 without it; setting the seedling entry to zero instead changes a predicted total by at most 43 stems.

Although the decadal projections largely reproduce each species' change, their sum over species does not match the observed change in the large-tree total (Table S20). Over two censuses the observed changes in the total range from 183 to 1,152 stems in absolute value, and for four of the five starting censuses the observed and projected changes have opposite signs. Summed over the five starts, the projected changes in the total are +168 stems at one census and +650 at two,

compared with +714 and +1,201 observed. These observed changes reflect variation shared across species, which a projection with the environment held at its mean cannot produce. With the shared year effect at the amplitude estimated from the transition data, the modeled total changes from census to census by amounts of the observed size (Results 3.3 of the main text; Section S8).

starting census	large trees	observed $\Delta_5 N$	predicted $\Delta_5 N$	observed $\Delta_{10} N$	predicted $\Delta_{10} N$
1982	20,952	+288	-166	+1,152	-208
1985	21,240	+864	-28	+1,004	+4
1990	22,104	+140	+107	-183	+257
1995	22,244	-323	+128	-578	+289
2000	21,921	-255	+126	-194	+308

Table S20: Predicted and observed change in the Barro Colorado large-tree total (stems ≥ 79.4 mm dbh, classes 5 to 8, the 84 modeled species) from each of the first five censuses, over one census interval and over two: the decadal change of Material and methods 2.6 of the main text summed over species, $\Delta_{10} N = \sum_s \Delta_{10} N_s$. The observed change is the large-tree total one or two censuses ahead minus the row's own. For the prediction the environment is held at its mean and the unobserved seedling entry of each stage vector is completed from the species' published fecundity row, as described in the text; with no seedling class at all the prediction changes by at most 27.1 stems. For most starting censuses the predicted change has the wrong sign as well as the wrong size, at both horizons, while for each species separately the same operators largely reproduce the changes (Fig. S4, Table S14).

S6 Test of neutral drift at Barro Colorado

Purpose: a test, on this plot alone, of whether the observed changes in large-tree abundance at Barro Colorado could be ecological drift among equivalent species.

The null hypothesis is neutral drift among equivalent species, Hubbell's model at the plot's own fitted parameters (Schwartz & O'Dwyer, companion paper). The community is of fixed size, and every tree has the same per capita rates of death and recruitment whatever its species. When a tree dies, the vacancy goes with probability $m = 0.026$ to an immigrant from a fixed source pool and otherwise to the offspring of a tree drawn at random from the plot. In that source pool we give each species its share of the plot at the starting census. Under this null, species identity is exchangeable given abundance: nothing else measured on a species at a census, whether size structure, transition rates or a model forecast, is informative about the sign or rank of its later change.

To express each species' observed change in units of neutral drift, let K_i be the number of large trees (stems ≥ 79.4 mm dbh) that died and were replaced on the plot in census interval i, taken from the tree tables, and p_i a species' share of the large trees. Under drift the species' number of large trees changes over the interval with variance $2K_i\, p_i(1-p_i)$, and summing from the starting census to 2015 gives the drift unit σ_s for its observed change ΔN_s,

$$\sigma_s = \Big[\sum_i 2K_i\, p_i(1-p_i)\Big]^{1/2}, \qquad z_s = \Delta N_s/\sigma_s. \tag{S2}$$

We repeat every calculation at two rates of replacement, the measured rate (K_i as observed) and the fitted rate (K_i multiplied by 4.51). The multiplier is the turnover ratio of the companion paper,

the factor by which the replacement rate fitted there to the plot's fluctuations exceeds the measured rate; it makes every drift unit 2.12 times larger.

The predictor is a species' demographic momentum in Keyfitz's sense (Caswell, 2001), computed from its size structure at a census and its own survival and growth rates from the 1982 to 1990 tree tables, with recruitment set to replace deaths. It is the factor by which the species' number of large trees would eventually change under those fixed rates, above one when the species holds a larger share of small stems than in its stable structure and below one when that share is smaller. Demographic momentum differs from the comparator labeled momentum in Table 2 of the main text, which extrapolates each species' previous change. Momentum is defined for the 75 of the 84 modeled species whose 1982 to 1990 tables yield a projection matrix and which have a large tree at the start. The statistic is the Spearman correlation S across those species between momentum and z_s. For the confidence quoted in the main text (first row of Table S21) we pair the momentum of the 1985 size structure with the change from 1990 to 2015. With this pairing every input to the momentum was censused by 1990, and no census interval contributes to both predictor and outcome.

We refer the observed correlation S to three null distributions of 200,000 draws each. In null (A) we simulate each species' large trees forward from the start under drift at the measured rate, holding the plot's totals at their censused values, and recompute S with the momenta unchanged; null (B) is the same at the fitted rate. In null (P) we reassign the momenta among species at random, which is valid under exchangeability for any rate of replacement and any form of the drift. Under each null, P is the fraction of draws, the observed value included, that reach the observed S; its smallest attainable value is thus one over the number of draws plus one. Drift at the measured rate, drift at the fitted rate and random reassignment of the momenta are three distinct hypotheses about how chance could produce the observed ranking, and we do not pool their draws. The bound $P < 10^{-5}$ is the strongest that 200,000 draws per null can support, and the 600,000 draws in total do not tighten it.

Table S21: Test of neutral drift at Barro Colorado alone. In each row the predictor is correlated across species with the change in large trees (stems ≥ 79.4 mm dbh) from the start census to 2015 in units of each species' neutral drift (Eq. (S2)); n, number of species with a projection matrix from the 1982 to 1990 tables and a large tree at the start; S, Spearman correlation. Columns A, B and P: how many of 200,000 draws reached S in forests simulated at the measured rate of replacement, at the fitted rate 4.51 times faster, and under random reassignment of the predictor among species; then the reassignment null's standard deviation and the largest 99.9th percentile and largest draw over the three nulls; $z_S = S\sqrt{n-1}$, a Gaussian approximation. The first row gives the confidence quoted in the main text. In the last two rows the census intervals used to compute the momentum also lie inside the outcome window (see text).

start	predictor	n	S	A	B	P	null s.d.	99.9th	largest	z_S
1990	momentum, 1985 structure	75	+0.667	0	0	0	0.116	+0.36	+0.52	5.7
1990	model forecast, 1990 structure	75	+0.728	0	0	0	0.116	+0.36	+0.52	6.3
1982	momentum, 1982 structure	75	+0.831	0	0	0	0.116	+0.36	+0.50	7.1
1985	momentum, 1985 structure	75	+0.775	0	0	0	0.116	+0.36	+0.53	6.7

Across the 75 species the 1985 momentum and the drift-standardized change from 1990 to 2015 have rank correlation $S = +0.67$, and none of the 600,000 draws from the three nulls reaches it, the largest being +0.52 (Table S21). Hence $P < 10^{-5}$ under each null, and neutral drift among equivalent species is rejected at this plot with more than 99.9% confidence. Replacing the

momentum by our model's forecast from 1990, each species' 1990 stage vector projected under its 1982 to 1990 rates and expressed in drift units, gives $S = +0.73$, which again no draw reaches. By 2015, 37 species lie outside their own 99% drift interval from 1990 at the measured rate and 14 at the fitted rate, compared with 0.75 expected under drift.

The observed correlation also exceeds every draw when the 1982 momentum is paired with the change from 1982 ($S = +0.83$) and the 1985 momentum with the change from 1985 ($S = +0.78$). In those two rows, however, the 1982 to 1990 transitions used for the momentum lie inside the outcome window. Under the null their sampling noise enters both predictor and outcome, a dependence that the three null distributions, which take the momenta as given, do not reproduce. If we divide the 1982 to 2015 change at 1990, its rank correlation with the 1982 momentum is +0.89 over the first window and +0.65 over 1990 to 2015. Because that first window also contributes 28% of the drift variance of the whole change, we base the stated confidence on the 1990 row.

The hypothesis rejected at this plot, at either rate of replacement, is drift among demographically equivalent individuals with immigration from a fixed pool. A neutral model in which all species share one set of size-dependent vital rates is a different hypothesis, and it is rejected too. Under it every stem of every species follows the community's per-interval survival and growth rates by size class (the pooled operator of the 84 species) from its species' observed 1990 size structure, entrants being allotted by each species' share of censused stems and immigrants at the same rate as above; the expected changes under this null rank the observed 1990 to 2015 changes at 0.22, and over 20,000 simulated forests the statistic S has mean +0.06 and standard deviation 0.02, no draw reaching +0.15 against the observed +0.67 ($P < 5 \times 10^{-5}$; the same with entrants allotted by large-tree share or with rates constant in time). Shared size-dependent rates acting on different size structures thus predict changes that did not occur, and the null is rejected more sharply than drift because its spread is narrower. What is not tested is a model whose rates are shared as a function of size relative to each species' stature, which is a life-history model in all but name. The companion paper (Schwartz & O'Dwyer) rejects drift on this plot by another statistic, the squared abundance change of every censused species against its exact neutral expectation, summarized by the turnover ratio used above, and we do not combine the two tests.

S7 Long-run composition in detail

Purpose: details behind Results 3.2: rank order and stem numbers by size class and over all censused stems, the two species with the largest discrepancies, the relation of local gain to observed abundance, and the sensitivity of the long-run results to growth-rate differences and to errors in the published stage rates.

Throughout, large trees are stems of at least 79.4 mm dbh, classes 5 to 8 of the published matrices (Jops et al., 2025). Stationary values are those of the closed form (Eq. (S1) of Section S2) at 73,729 immigrant seeds per interval or, where stated, means over the stationary simulations. Under the regional pool we computed the rank correlation of stationary with census-mean abundance within one large-tree class at a time, over the species whose matrix contains that class (Table S22, second part). The correlation is highest in class 5, the first large-tree class, which holds 63% of the observed large trees, and lower in the three larger classes. We restrict each correlation to species whose matrix contains the class because a species whose published matrix lacks class 7 or 8 holds no stems there in the model, and none in the census under the same class convention, which assigns a species' largest observed stems to its top fitted class. For the 43 of the 84 species that lack class 7 and the 65 that lack class 8, agreement in those classes would otherwise be built in.

Between species, the regional pool's stationary log abundances have a standard deviation of

1.14, against 0.76 observed ($\log_{10}$ units), and the identity-line R^2 remains negative when *Xylopia macrantha* is excluded (−0.649). The capacity constrains only the total number of censused stems and leaves their distribution among size classes free. Under the regional pool 27.95% of censused stems stand in the large-tree classes, against 10.81% in the census (Table S22, first part; Fig. S2a). This excess does not arise within species, because at the stationary state every species keeps its stable stage structure (Eq. (S1)), which places a share of its own censused stems in the large-tree classes close to the observed share (Table S22, last part). At the observed species composition these structures would put 13.03% of the stand in those classes. The excess comes instead from species composition: under the regional pool the stationary stand is dominated by species whose stable structures are weighted toward the large-tree classes, chiefly *X. macrantha*. Under the compensated scheme, which shifts composition away from such species, the large-tree share is 13.52%.

When we compare predicted with observed abundance on all censused stems (at least 10 mm dbh) instead of on large trees, the rank correlation falls to 0.261 under the regional pool and to 0.346 under the compensated scheme. The identity-line R^2 is strongly negative under both schemes (Table S22, third part). The stems in this comparison are mainly saplings, which are 89.19% of censused stems and belong mostly to understory species that rarely or never reach 10 cm dbh. The observed spread of log abundance across species is also narrower than for large trees (standard deviation 0.45 against 0.76 orders of magnitude), whereas each scheme's predicted spread exceeds it by a larger factor than for large trees. The regional shares p_s themselves have a rank correlation of only 0.24 with all-stem abundance. Class by class the failure among saplings is one of composition rather than level. Within censused classes 2, 3 and 4 the rank correlation of stationary with census abundance is 0.23, 0.34 and 0.45 under the compensated scheme, whose stationary total in each of those classes is 0.99, 0.91 and 1.03 times the census, and 0.15, 0.26 and 0.43 under the regional pool (totals 0.67, 0.78 and 1.22); over the three sapling classes together the correlations are 0.28 and 0.22 with identity-line R^2 of −2.08 and −6.21. Over the 79 species whose matrices reach the large-tree classes, the large-tree register reads 0.67 ($R^2 = +0.12$) under the compensated scheme and 0.60 (−1.95) under the pool, against 0.727 and 0.666 over all 84. The error enters with the recruits rather than in the matrices' transfer between classes: the observed seedling composition passed through the published matrices orders the census saplings at 0.73 and the large trees at 0.71, although at that composition the published seedling promotion rates would supply only about a third of the observed sapling stand (Section S9).

The two largest discrepancies under the regional pool are that *X. macrantha* is the predicted dominant and that the observed dominant, *Faramea occidentalis*, ranks far down the predicted order (Fig. 4a of the main text). By regional share alone *X. macrantha* would rank 11 of 84; it is nevertheless the predicted dominant because, under its published rates, its ratio Ω_s/μ_s of large trees per seedling to seedling mortality lies far above the community median (Table S23). Replacing or removing *X. macrantha*'s matrix does not, however, repair the regional pool's absolute abundances (counterfactuals in Table S23): without that species the large-tree total is still 1.20 times the observed, and another long-lived species becomes the predicted dominant.

Among the 79 species whose matrices reach the large-tree classes, the ratio Ω_s/μ_s, to which the local gain G_s is proportional (Section S2), has a standard deviation of 0.84 in $\log_{10}$ units, 92% of that variance coming from Ω_s, the large trees per seedling, and the rest from the seedling mortality μ_s. By Eq. (S1) the gain that would reproduce a species' observed large-tree abundance under the regional pool is proportional to that abundance divided by its regional share. Across these species this required gain decreases with the gain computed from the matrix (log-log slope −0.20, 90% CI −0.32 to −0.07), whereas if the computed gains acted as modeled the slope would be one. A large local gain under shared recruitment thus does not make a species common at Barro Colorado. The slope lies 16 standard errors below the value of one the regional pool implies (in a bootstrap over

Stems by censused class relative to the census mean, and the share standing in classes 5 to 8	class 2	class 3	class 4	class 5	class 6	class 7	class 8	% in 5 to 8
census mean (stems)	80,071	66,597	32,360	13,727	5,509	2,080	377	10.81
regional pool / census	0.67	0.78	1.22	2.30	3.93	1.14	1.33	27.95
compensated scheme / census	0.99	0.91	1.03	1.14	1.59	1.04	1.62	13.52

Rank correlation with the census within each large-tree class, regional pool (simulations)	class 5	class 6	class 7	class 8
species with the class in their matrix	0.59	0.52	0.39	0.45
(number of such species)	(79)	(62)	(40)	(19)

The species comparison made on all censused stems, classes 2 to 8	ρ, classes 2 to 8			R^2 of log	sd ratio,	ρ, classes
weighting	full	rare	common	abundance	pred./obs.	5 to 8
regional pool, $w_s \propto p_s$	0.261	+0.14	−0.14	−7.80	2.41	0.666
compensated scheme, $w_s \propto p_s/G_s$	0.346	+0.11	+0.23	−2.45	1.78	0.727

Share of each species' censused stems in classes 5 to 8, predicted against observed (79 species)	
median predicted share	0.144
median observed share	0.139
predicted over observed, quartiles across species	0.62 to 1.50
rank correlation, predicted with observed	0.77

Table S22: Size classes at the stationary state under the two immigration schemes, 84 species at Barro Colorado; values from Eq. (S1) at 73,729 immigrant seeds per interval except the second part (stationary simulations). First part: census-mean stems per censused diameter class (class 2 begins at 10 mm, class 5 at 79.4 mm dbh) and each scheme's stationary number as a ratio to it; the capacity fixes the classes 2 to 8 total at the observed value, and an excess among large trees is therefore offset by a deficit among saplings. Second part: Spearman correlation of stationary with census-mean abundance within each large-tree class under the regional pool, over the species whose published matrix contains the class. Third part: the comparison of Results 3.2 repeated on all censused stems (rank correlation over all species and over the rarer and commoner halves, identity-line R^2 on log abundance, ratio of predicted to observed standard deviation of log abundance across species), with each scheme's large-tree rank correlation for reference. Last part: each species' predicted share of its own censused stems in classes 5 to 8, a property of its published matrix alone, against the observed share, over species whose matrices reach those classes.

species the probability of a non-negative slope is 0.0005); with the regional share's own exponent left free the gain's exponent is $+0.03 \pm 0.06$ (the share's 0.52); and the rejection does not rest on the least certain entries, since with seedling survival and promotion set equal across species the slope is -0.12 (90% CI -0.27 to $+0.04$), still 13 standard errors below one, while at growth rate one the fecundity row cannot enter G_s at all. The gain does, however, mark the direction of recent change: between 1985 and 2015 the large-tree shares of 53 of the 79 species moved toward the pool's stationary composition $p_s G_s$ (two-sided binomial $P = 0.003$), the movement having a rank correlation of $+0.58$ (90% CI $+0.44$ to $+0.69$) with $\log G_s$, at a rate that would close the gap only over some three centuries per e-folding; no such movement toward the compensated scheme's composition is detectable. The census therefore cannot separate arrivals that offset the gain from a stand still centuries from the uncompensated state, and the gains and the movement come from rates of the same decades. Two measurements would separate them: the composition of arriving seed relative to regional share, which should fall with the local gain under compensation, and the stationary composition of a second forest with published matrices, which should rise with it under the pool.

	X. macrantha	*F. occidentalis*	median species
large trees, census mean	211	3,391	
large trees, regional pool (rank)	30,387 (1)	217 (23)	
large trees, compensated scheme (rank)	706 (11)	(10)	
stems in classes 5 and 6, regional pool (simulations)	16,581, 13,789		
stems in classes 5 and 6, census	167, 44		
Ω_s/μ_s, large trees per seedling over seedling mortality	2.85	0.020	0.022
share of own censused stems in classes 5 to 8, predicted (observed)	0.56 (0.18)	0.21 (0.13)	0.144 (0.139)
Counterfactuals under the regional pool, Eq. (S1)*, each growth rate held at one*			
X. macrantha with the across-species median seedling column: 3,518 large trees, rank 4; $\rho = 0.667$, $R^2 = -0.53$; large-tree total 1.74 × observed			
X. macrantha with *F. occidentalis*'s whole matrix and its own regional share: 290 large trees; $\rho = 0.668$, $R^2 = -0.49$; total 1.63 × observed			
X. macrantha left out of the comparison: the other species give $\rho = 0.669$, $R^2 = -0.60$; total 1.20 × observed			

Table S23: The two species with the largest discrepancies under the regional pool at Barro Colorado, *Xylopia macrantha* (predicted dominant) and *Faramea occidentalis* (observed dominant), against the median of the 84 species where one is defined. Large trees are stems of at least 79.4 mm dbh (classes 5 to 8); stationary values from Eq. (S1) at 73,729 immigrant seeds per interval, except *F. occidentalis*'s large trees under the regional pool and the class 5 and 6 counts, which are means over the stationary simulations. The base values of the regional pool for the counterfactuals are $\rho = 0.666$, $R^2 = -0.693$ and a large-tree total 2.59 times the observed 21,693.

Unlike the rank order, absolute abundances and the identity of the leading species depend on the normalization of every published matrix to growth rate one. Under that normalization each species' growth rate, once the stand is full, is 0.0003 to 0.03 below replacement per interval, and its stationary abundance is its immigrant supply divided by that gap (Section S2). When we raise one species' fecundity until its growth rate is 1.005, leave the others at one and re-solve the closed form under the regional pool, its abundance increases by a median factor of 3.3, and in six of the 84 cases

the raised species displaces *X. macrantha* as the predicted dominant. When instead we perturb one species' survival in every class by the same half percent and rescale its fecundity to growth rate one, the kind of error the published estimates may contain, its abundance changes by a median factor of only 1.11 and the predicted dominant is the same in every case. Across all these single-species perturbations the rank correlation remains between 0.64 and 0.69. The absolute predictions of Results 3.2 therefore hold for species of equal long-run fitness and would change under growth-rate differences larger than those gaps below replacement. The censuses themselves do not certify equality at that precision. We redrew every transition probability of classes 2 to 8 of every species from a multinomial at its census sample size (1.4 million stem-intervals in all), kept the fecundity rows and the seedling column, and re-solved the closed form 500 times under each scheme without renormalizing the matrices to growth rate one. Species' long-run growth rates then scatter with a median standard deviation of 0.0065, comparable to the median gap below replacement under the regional pool (0.007) and larger than that under the compensated scheme (0.001), and about 43 of the 84 species are self-sustaining in a typical draw. The predicted dominant then differs from the one of record in 89% of draws under the pool and 98% under compensation, the identity-line R^2 on log abundance is negative in every draw (median -2.9 under compensation, against $+0.495$ of record), and the large-tree total ranges several-fold either way; the rank correlation with the census, by contrast, has a median of 0.64 (5th to 95th percentile 0.59–0.68) under the pool and 0.70 (0.65–0.74) under compensation, against 0.666 and 0.727 of record, because regional share still orders the immigration-maintained species (0.59 and 0.64 when the seedling column is redrawn as well). The decadal forecasts are untouched, since the normalization enters them only through fecundity, which does not reach the large-tree classes in two steps (Section S5). The rank order and the forecasts are thus the model's tested predictions; the total, the R^2, the named dominant and the abundance distribution are its consequences at exactly equal long-run fitness.

Because the closed form gives every long-run statistic at any set of rates and the decadal forecast uses the matrices directly, we obtained the effect of an error in the published stage rates without further simulation. We changed the rates by 20% up and down in five groups, one at a time (the four row groups of Table S24 and the fecundity row), for every species at once, for *X. macrantha* alone and for *F. occidentalis* alone, rescaling each affected fecundity row to growth rate one. That rescaling cancels a uniform change of the fecundity row, so that the fecundity group acts only in the unnormalized scenarios below and 24 renormalized scenarios remain. In each scenario we evaluated the long-run statistics under the regional pool, under the compensated scheme with its gains G_s recomputed from the perturbed rates and under a mis-specified compensated scheme whose arrivals are weighted by the unperturbed gains, and the out-of-sample decadal forecast of Section S5.

The decadal forecast uses only the survival and promotion entries of the censused classes (Section S5). A 20% error in sapling promotion or in adult mortality applied to every species alters the forecast's rank correlation by at most 0.06 and leaves its margin over the forecast from each species' own trend at least $+0.11$. Under the regional pool the long-run order largely follows that of the regional shares, and in every renormalized scenario the rank correlation lies between 0.63 and 0.68, R^2 is negative and *X. macrantha* remains the predicted dominant. The rate changes chiefly alter that species' predicted abundance (Table S24, last columns) and, through it, how far below zero R^2 lies (-0.58 to -1.40). Under the compensated scheme with the gains recomputed, the stationary composition is the regional one in every scenario, and the statistics change by at most 0.002 in rank correlation and 0.07 in R^2. When that scheme is mis-specified, a one-species error scales the species' abundance by the same factor as its gain and changes the community's rank correlation and R^2 by at most 0.005 and 0.03. An error common to all species lowers R^2 as far as $+0.32$ only when every species' seedling survival lies above its estimate at once.

The bounds above hold only under the normalization to growth rate one, and without the

rescaling of fecundity the same changes shift the growth rates away from one. A species whose growth rate is raised above one then sustains itself without immigration and takes over the censused stand (15 of 30 scenarios). Where the affected species' growth rates fall below one (15 scenarios) the stand is maintained by immigration alone, and in 4 of these it no longer fills to the capacity at the arrival rate used here. With every growth rate held at one, the uncertain seedling and fecundity entries thus affect mainly one quantity, the size of the regional pool's over-prediction of absolute abundances reported in Results 3.2. The forecast, the rank order and the compensated abundances do not depend on them materially.

rates changed by ±20%	species	regional pool		compensated		compensated, mis-specified		forecast	*X. macrantha*, pool large trees (base 30,387)	
		$\lvert\Delta\rho\rvert$	$\lvert\Delta R^2\rvert$	$\lvert\Delta\rho\rvert$	$\lvert\Delta R^2\rvert$	$\lvert\Delta\rho\rvert$	$\lvert\Delta R^2\rvert$	$\lvert\Delta\rho\rvert$	low	high
class-1 survival	all	0.034	0.71	0.001	0.01	0.031	0.18	–	23,700	27,300
	Xylopia	0.000	0.31	0.001	0.00	0.003	0.03	–	16,900	65,000
	Faramea	0.005	0.01	0.001	0.01	0.005	0.01	–	30,300	30,400
class-1 promotion	all	0.000	0.01	0.001	0.01	0.001	0.01	–	30,400	30,400
	Xylopia	0.000	0.03	0.001	0.00	0.001	0.01	–	25,700	34,600
	Faramea	0.004	0.01	0.002	0.01	0.004	0.01	–	30,400	30,400
sapling promotion, classes 2 to 4	all	0.002	0.17	0.001	0.07	0.012	0.07	0.051	27,100	32,800
	Xylopia	0.000	0.01	0.002	0.01	0.001	0.01	0.001	26,600	33,300
	Faramea	0.005	0.01	0.001	0.01	0.003	0.01	0.000	30,400	30,400
adult mortality, classes 5 to 8	all	0.006	0.05	0.002	0.02	0.003	0.02	0.055	27,200	34,500
	Xylopia	0.000	0.02	0.001	0.01	0.001	0.01	0.001	26,700	35,500
	Faramea	0.002	0.01	0.001	0.01	0.002	0.01	0.000	30,400	30,400

Table S24: Sensitivity of four long-run and forecast statistics to an error of 20% in one group of the published stage rates, each species' fecundity then rescaled to growth rate one. Entries are the larger of the two changes (rates raised, rates lowered) from the base values: $\rho = 0.666$, $R^2 = -0.693$ for the regional pool; 0.727, +0.495 for the compensated scheme (arrivals $w_s \propto p_s/G_s$ with G_s recomputed from the changed rates); the same with G_s left at the unchanged rates (mis-specified); and the out-of-sample rank correlation of the decadal forecast (Section S5), the larger change over one and two censuses ahead from base values 0.71 and 0.68; a dash marks the two class-1 groups, whose rates do not enter the forecast because it uses only the survival and growth frequencies of the censused classes. Last columns: *Xylopia macrantha*'s large trees under the regional pool with the rates lowered and raised, to the nearest hundred. All values from Eq. (S1) at 73,729 immigrant seeds per interval and the perturbed matrices; no simulation is involved.

S8 Fluctuation predictions

Purpose: the criteria and tolerances for comparing the year effect's consequences with the censuses, and the outcomes under both immigration schemes, with and without species-level mortality variation.

Because the form of the year effect (the transitions it multiplies) and its amplitude σ_g are estimated from the transition data and its memory τ_g is assumed (Section S3), its consequences for large-tree numbers (stems ≥ 79.4 mm dbh) involve no constant fitted to an abundance or a count of stems; they are compared here with the same seven intervals from which σ_g was estimated, and later censuses will test them out of sample. We ran the model to its stationary state under the regional pool ($w_s = p_s$) and the compensated scheme ($w_s \propto p_s/G_s$), 36 replicates of each setting, all with $\tau_g = 18$ yr (Section S2). We set the amplitude at $\sigma_g = 0.32$, at either end of its 90% confidence interval (0.17 and 0.84) and at zero (the factor off). We report each model statistic as its median

and range over consecutive non-overlapping windows of seven steps, the length of the census series, spanning the last 2,000 steps of each replicate. We did not vary the memory separately: over the assumed range of τ_g, 5 to 40 yr, the estimated amplitude remains inside that confidence interval.

For recruitment the statistic is the coefficient of variation (CV) of recruitment into class 5, the first large-tree class, over seven censuses (criterion 0.14 to 0.22 about the measured 0.18; Section S3). For the large-tree total the statistics are its median absolute five-year relative change (criterion within a factor of two of the census's 0.0127) and its thirty-year range ratio, largest over smallest in a window (below 1.10; census 1.062). The standing total should stay within 20% of its value with the factor off, and the rank correlation of the stationary large-tree abundances with the census means should change by no more than five hundredths. The stationary seedling total should lie within a factor of three of the seedling-census mean (Section S9).

Within each of ten bands of initial large-tree abundance on a doubling scale we compute the mean squared five-year change of a species' large-tree count. In the two most abundant bands (species of 256 or more large trees) the census value should lie within a factor of two of the model's. We also divide the census value in each band by the model's value with the factor off. In at least seven bands this ratio should fall inside the range, over windows, of the model's own ratio of its value with the factor on to its value with the factor off. Finally, we measure, identically in model and census, how much of each band's five-year change occurs in proportion to the change of the whole large-tree total. A species' coherent change is its share of the large-tree total times that total's change, and a band's coherent share is the sum over its species of the squares of these coherent changes, divided by the sum of their squared five-year changes.

With the factor off, the median absolute five-year relative change of the large-tree total under the regional pool is 0.0011, against 0.0127 in the census, smaller by a factor of 12 (Results 3.3 of the main text). At the estimated amplitude the regional pool meets every criterion on recruitment, on the large-tree total and on the composition, and the compensated scheme all but two of them (Table S25). Its range ratio slightly exceeds the limit, and its standing total settles at 0.68 of the value without the factor, a thinning for which we give two interpretations in Results 3.3 of the main text. Under the regional pool the modeled large-tree total agrees with the census only in its relative fluctuations, because its standing value without the factor is 2.59 times the census mean (Section S2). At the upper end of the amplitude's confidence interval the census's mean squared change lies inside the model's range in more bands than at the estimated amplitude. However, the recruitment CV is then three times the measured value, the species of both top bands fluctuate at least twice as much as observed, and under the compensated scheme the rank correlation changes by more than five hundredths. The variability of recruitment and the five-year changes of the commonest species thus exclude amplitudes in the upper part of the confidence interval that the transition data give for σ_g.

The census's mean squared change lies within the factor of two of the model's in both top bands under the regional pool and in the lower of them under the compensated scheme. It lies inside the model's range in 5 of the ten bands under the regional pool and 6 under the compensated scheme, fewer than the seven required. The shortfall is at intermediate abundance: for species of 16 to 255 large trees the census's mean squared change remains, with the factor on, 2.4 to 3.3 times the model's under the regional pool and similarly under the compensated scheme (Table S25). In the top band the model's coherent share exceeds the census's because most of the observed variation there is each species' own steady trend (78% of those species' summed squared five-year change), which the projection matrices account for (Section S5). After detrending, the abundant species in the census share 29% of their remaining squared change with the total (95% CI 12–57%). Simulations without the factor give 2%, and we attribute the difference to environmental variation common to all species, as represented by the year effect.

Two terms measured in the transition data but omitted from the model could supply the census-to-census variability the model lacks at intermediate abundance (Section S3): species-specific deviations of sapling growth from the common year effect, and each species' own year-to-year variation in mortality. When the second is added to the simulations at its measured size (lower rows of Table S25), every species fluctuates more between censuses, and the increase is largest for abundant species. The census's mean squared change then lies within the factor of two of the model's in both top bands under either scheme, and inside the model's range in 7 of the bands under the regional pool and 7 under the compensated scheme. Every other criterion in the table is met under both schemes as well, and under the compensated scheme the stand no longer thins. For species of intermediate abundance the census's mean squared change is still 2.1 to 2.7 times the model's under the regional pool, leaving the species-specific growth deviations, not simulated here, as the remaining candidate.

The year effect also widens the prediction intervals of the species-level abundance forecasts. In the neutral-drift test (Section S6) each species' 1990 stage vector is projected to 2015 under its matrix from the 1982–1990 tree tables. We repeated those projections for the 75 species of that test, 2,000 simulations each, with the year effect at its estimated amplitude and common to all species. The central 90% of the simulated changes then covers 60% of the observed changes, compared with 27% with demographic noise alone and 96% at the upper end of the amplitude's confidence interval. At the estimated amplitude the year effect thus accounts for part of the spread the forecasts lack, without bringing their coverage to the nominal 90%. These forecasts are also biased upward, mainly because their matrices come from the two intervals of strongest sapling growth; sapling promotion rates over the forecast window were 0.58 to 0.79 of their 1982–1990 values. Rescaling the class 2 to 4 promotions by those ratios, interval by interval, removes 92% of the mean bias with the rank correlation little changed, but uses conditions realized within the forecast window and so explains the bias only after the fact.

Table S25: Consequences of the year effect estimated from the transition data, under the two immigration schemes, at the estimated amplitude ($\sigma_g = 0.32$) and at the two ends of its 90% confidence interval, all at $\tau_g = 18$ yr (model; 36 replicates per row, statistics over seven-step windows of their last 2,000 steps, each compared with the same scheme with the factor off). Columns: the coefficient of variation (CV) of recruitment into class 5, the first large-tree class, over seven censuses, median over windows (fifth to ninety-fifth percentile at $\sigma_g = 0.32$), tolerance 0.14–0.22 about the measured 0.18; the number of the ten abundance bands in which the census's mean squared five-year change lies inside the model's range over windows, seven required; the median absolute five-year relative change of the large-tree total (stems $\geq$ 79.4 mm dbh), tolerance a factor of two about the census; the thirty-year range ratio of that total, tolerance below 1.10; the standing large-tree total over its value with the factor off, tolerance 20% either side of one; the census's mean squared five-year change over the model's with the factor on in the two most abundant bands (species of 256 or more large trees, the lower band first), tolerance a factor of two, and its range over the four bands of species with 16 to 255 large trees (mid bands); the coherent share (%) of the most abundant band's squared five-year change with the factor on, the census value being that before detrending (29% after; see text); the change that the factor makes in the rank correlation of the stationary abundances with the census means, tolerance five hundredths. In the upper rows the range ratio, the mid-band ratio and the coherent share are given at the estimated amplitude only. In the lower rows (marked mortality), the same statistics with each species' own year-to-year variation in mortality added: a mean-one lognormal factor on the five-year mortality hazard, drawn independently for every species at every step, with the standard deviations measured in the transition data for small and for large stems (Section S3), alone (year effect off) and together with the year effect at $\sigma_g = 0.32$ (36 replicates each, compared with the same scheme with both off). Over replicates the median relative change of the total with both terms ranges from 0.0157 to 0.0178 under the regional pool (year effect alone 0.0145 to 0.0167) and from 0.0138 to 0.0227 under the compensated scheme (0.0181 to 0.0216). With the year effect's memory given to the mortality factor as well, that median is 0.0175 under the regional pool and 0.0236 under the compensated scheme. Last row, the census values.

scheme	σ_g	CV, class 5	inside	total	range	stand	top bands	mid bands	coherent	$\Delta\rho$
regional pool	0.17	0.11	2	0.009	–	1.00	3.19, 5.17	–	–	0.000
	0.32	0.21 (0.10–0.45)	5	0.016	1.088	0.97	1.67, 1.88	2.4–3.3	35	+0.003
	0.84	0.55	9	0.039	–	0.82	0.39, 0.50	–	–	−0.001
compensated	0.17	0.11	3	0.011	–	0.86	2.75, 5.47	–	–	−0.009
	0.32	0.21 (0.10–0.45)	6	0.020	1.109	0.68	1.40, 2.52	2.6–3.1	58	−0.025
	0.84	0.56	9	0.043	–	0.56	0.36, 0.47	–	–	−0.073
regional pool, mortality	off	0.02	2	0.0052	1.022	0.99	2.55, 2.99	–	6	+0.001
	0.32	0.21	7	0.0169	1.093	0.96	1.23, 1.23	2.1–2.7	26	+0.004
compensated, mortality	off	0.03	5	0.0060	1.026	1.14	1.63, 2.00	–	6	−0.040
	0.32	0.21	7	0.0172	1.096	1.01	0.95, 1.20	1.9–2.7	21	−0.038
census		0.18		0.0127	1.062				7.5	

S9 The seedling pool by species

Purpose: the model's stationary seedling class compared with the Barro Colorado seedling census under both immigration schemes, in total and by species, and the rank correlations of the observed seedling composition with regional and with large-tree abundance.

The rates of the seedling class (class 1) in the published matrices (Jops et al., 2025) were estimated from the Barro Colorado seedling census (Comita et al., 2023a,b), which covers every free-standing woody stem at least 20 cm tall and under 1 cm in diameter. Such stems were enumerated in 19,313 plots of 1 m^2 at fourteen censuses from 2001 to 2018. Scaled to the 50-ha plot, the counts of the modeled species averaged 1.13 million stems over those censuses and ranged from 0.93 to 1.39 million at individual censuses. These species made up 67 to 75% of the live seedlings of the 413 species found in the plots. Stems shorter than 20 cm are not censused but belong to class 1 (every stem under 1 cm dbh), and these counts therefore understate class-1 abundance.

In Table S26 we compare the seedling census with the model's stationary class 1, computed from the closed form of Section S2 without the year effect (the shared year-to-year factor on sapling growth). Under the regional pool, with arrivals in proportion to regional relative abundance, the seedling class holds 1.25 million stems, 1.10 times the census mean and 0.89 to 1.35 times the individual censuses, within the factor-of-three criterion of Section S8. Under the compensated scheme, arrivals are in proportion to regional abundance divided by the local gain G_s, the number of large trees (stems of at least 79.4 mm dbh) standing per arriving seed. With this weighting the seedling class holds 5.44 million stems, 4.8 times the census mean and more than three times every individual census.

By Eq. (S1) each arriving seed of species s leaves $1/(d\mu_s)$ seedlings standing, where μ_s is the seedling mortality per step and d the fraction of candidate recruits not admitted, so that every species' stationary seedling class is proportional to $1/d$. Because the compensated scheme weights arrivals toward species whose seeds yield few large trees, the censused stand is held at capacity only if a larger share of candidate recruits is admitted each step, with $d = 0.026$ in place of 0.153 under the regional pool. This smaller d accounts for the compensated scheme's larger seedling class. That increase is less than the factor $1/d$ alone would give, because the same weighting favors species with high seedling mortality μ_s, whose seeds leave fewer seedlings standing. The arrival rate cannot remove the excess: the capacity fixes $M(1-d)/d$, so a lower $M = \hat{m}B$ lowers d with it and takes away only the immigrant seeds' own seedlings, the fraction $d = 0.026$ of the class (Section S2). Equivalently, the seed production of the standing trees that holds the compensated scheme's stand at capacity is several times the regional pool's, and it, not the nominal B, sets the seedling class. The excess therefore counts against the compensated scheme's stationary stand at the seedling stage, where the model has no density dependence, and not against its censused predictions, none of which involve the seedling class's size.

For individual species neither scheme reproduces the observed seedling abundances. The regional pool places 28 of the 84 species within a factor of three of their observed mean, and the compensated scheme places 36. Across species we find that the rank correlation of predicted with observed seedling abundance is 0.01 under the regional pool and 0.26 under the compensated scheme. The observed seedling composition resembles neither the regional composition nor that of the plot's large trees. Regional relative abundance, whose rank correlation with observed large-tree abundance is 0.72, correlates with the observed seedling composition at −0.02, and the observed seedling and large-tree compositions correlate at +0.23. Because the model's seedling class has no density dependence of its own and only supplies recruits to the censused stand, we should perhaps not expect either immigration scheme to reproduce the observed seedling composition, which may

	seedling census 2001–2018, mean (range over 14 censuses)	regional pool, stationary class 1 (ratio to the census mean)	compensated scheme, stationary class 1 (ratio to the census mean)
84 species, millions on the plot	1.13 (0.93–1.39)	1.25 (1.10; by census 0.89–1.35)	5.44 (4.8; by census 3.9–5.9)
84 species, per m^2	1.9–2.8	2.5	10.9
84 species, share of all live seedlings (413 species found), %	67–75		
fraction of candidate recruits not admitted, d		0.153	0.026
Xylopia macrantha, stems on the plot	4,300	55,200 (12.7)	1,100 (0.26)
Faramea occidentalis, stems on the plot	97,800	23,100 (0.24)	66,100 (0.68)
species within a factor of three of their census mean		28 of 84	36 of 84
species above, below that band		28, 28	38, 10
rank correlation with observed seedling abundance, across species		0.01	0.26

Table S26: The model's stationary seedling class (class 1) compared with the Barro Colorado seedling census under the two immigration schemes, without the year effect. Observed: live free-standing woody stems at least 20 cm tall and under 1 cm dbh of the 84 modeled species in the 19,313 plots of 1 m^2 at each of 14 censuses, scaled to 50 ha; species rows are means over the censuses to the nearest hundred. Model: the closed-form stationary state of Section S2 under the regional pool (arrivals in proportion to regional relative abundance) and under the compensated scheme (arrivals in proportion to regional abundance divided by local gain). Ratios are model over the observed 14-census mean; "by census" gives the range of that ratio over the individual censuses. Stems shorter than 20 cm are not censused, so the observed values are lower bounds. The regional pool matches the observed total within its census-to-census range and the compensated scheme exceeds it by more than a factor of three. Neither scheme reproduces the rank order of the species' seedling abundances.

instead reflect conspecific density dependence in seedling survival (Results 3.4).

S10 Data sources and code archive

Purpose: the repository, DOI and license of the public datasets we use, and the archive of our code and derived tables.

All data we analyze are public. The 1982–2015 tree censuses of the 50-ha plot on Barro Colorado Island (BCI) are the complete-plot release on Dryad (Condit et al., 2019, doi:10.15146/5xcp-0d46, CC0); Condit (1998) and Hubbell et al. (1999) describe the plot and its census methods. The BCI seedling and small-sapling census is the Dryad release of Comita et al. (2023a, doi:10.5061/dryad.05qfttf85, CC0), described by Comita et al. (2023b). Each species' stage-structured projection matrix is from the Zenodo archive of Jops and O'Dwyer (2024, doi:10.5281/zenodo.14563651, CC BY 4.0) that accompanies Jops et al. (2025). We build the regional species pool from the Panama plot-network tables of Condit et al. (2016, Smithsonian, doi:10.5479/data.stri.2016.0622, CC BY 4.0), whose current archive is on Dryad (Condit

et al., 2022, doi:10.15146/mdpr-pm59, CC0). For the exploratory comparisons of Section S3 we use further public data. Species' demographic rates in high and low light and their maximum heights are from the Dryad tables of Rüger et al. (2022, doi:10.7291/D16M46, CC0). The sapling growth and mortality indices, wood density, seed mass and stature are from the figshare data supplement (doi:10.6084/m9.figshare.3550359.v1, CC0) of Wright et al. (2010); for species missing from that supplement we use the maximum height listed by Condit et al. (2020, Dryad, doi:10.15146/R3M97W, CC0) as the stature. Each species' dry-season moisture association is the index (Smithsonian, doi:10.5479/data.bci.20130204, CC BY 4.0) published with Condit et al. (2013). The BCI climate series (monthly rainfall, solar radiation and air temperature, and dry-season dates) are the Smithsonian Tropical Research Institute dataset of Paton (2024, doi:10.25573/data.10059455.v46, CC BY 4.0), and the Oceanic Niño Index is the public-domain series of the NOAA Climate Prediction Center. Our code and the derived tables it produces are archived, with a README, at [repository and DOI to be added before submission]. In particular, the archive contains a derived table of the projection matrices used in the out-of-sample forecasts, which we estimated from the 1982–1995 censuses only, and the code that computes those matrices from the census release. The archived code reproduces every number, table and figure in the main text and in this Supporting Information from the public inputs listed above.

References

Caswell, H. (2001). *Matrix Population Models: Construction, Analysis, and Interpretation*. Sinauer Associates, Sunderland, MA, second edition.

Comita, L., Aguilar, S., Hubbell, S., and Pérez, R. (2023a). Long-term seedling and small sapling census data from the Barro Colorado Island 50 ha Forest Dynamics Plot, Panama. Dryad, dataset.

Comita, L. S., Aguilar, S., Hubbell, S. P., and Pérez, R. (2023b). Long-term seedling and small sapling census data from the Barro Colorado Island 50 ha Forest Dynamics Plot, Panama. *Ecology*, 104(9):e4140.

Condit, R. (1998). *Tropical Forest Census Plots: Methods and Results from Barro Colorado Island, Panama and a Comparison with Other Plots*. Springer, Berlin, Heidelberg.

Condit, R., Engelbrecht, B. M. J., Pino, D., Pérez, R., and Turner, B. L. (2013). Species distributions in response to individual soil nutrients and seasonal drought across a community of tropical trees. *Proceedings of the National Academy of Sciences*, 110(13):5064–5068.

Condit, R., Pérez, R., and Aguilar, S. (2020). Complete tree species of Panama. Dryad dataset, version 2, doi:10.15146/R3M97W.

Condit, R., Pérez, R., Aguilar, S., and Lao, S. (2016). Data from tree censuses and inventories in Panama. Smithsonian Institution, dataset (plot tables of January 2013).

Condit, R., Pérez, R., Aguilar, S., and Lao, S. (2022). Census data from 65 tree plots in Panama, 1994–2015. Dryad dataset.

Condit, R., Pérez, R., Aguilar, S., Lao, S., Foster, R., and Hubbell, S. (2019). Complete data from the barro colorado 50-ha plot: 423617 trees, 35 years. Dryad dataset, doi:10.15146/5xcp-0d46.

Hastings, A. (1980). Disturbance, coexistence, history, and competition for space. *Theoretical Population Biology*, 18(3):363–373.

Hubbell, S. P., Foster, R. B., O'Brien, S. T., Harms, K. E., Condit, R., Wechsler, B., Wright, S. J., and de Lao, S. L. (1999). Light-gap disturbances, recruitment limitation, and tree diversity in a neotropical forest. *Science*, 283(5401):554–557.

Jops, K., Dalling, J. W., and O'Dwyer, J. P. (2025). Life history is a key driver of temporal fluctuations in tropical tree abundances. *Proceedings of the National Academy of Sciences*, 122(4):e2422348122.

Jops, K. and O'Dwyer, J. (2024). Space-limited matrix community model - Barro Colorado Island tree species. Zenodo, dataset.

Jops, K. and O'Dwyer, J. P. (2023). Life history complementarity and the maintenance of biodiversity. *Nature*, 618(7967):986–991.

Levins, R. and Culver, D. (1971). Regional coexistence of species and competition between rare species. *Proceedings of the National Academy of Sciences*, 68(6):1246–1248.

Paton, S. (2024). Monthly summary_BCI, horizontal. Smithsonian Tropical Research Institute dataset, version 46, doi:10.25573/data.10059455.v46.

Rüger, N., Hubbell, S. P., and Condit, R. (2022). Demographic response to light environment of all species in the Barro Colorado plot: recruitment, growth, and mortality. Dryad dataset, doi:10.7291/D16M46.

Tilman, D. (1994). Competition and biodiversity in spatially structured habitats. *Ecology*, 75(1):2–16.

Wright, S. J., Kitajima, K., Kraft, N. J. B., Reich, P. B., Wright, I. J., Bunker, D. E., Condit, R., Dalling, J. W., Davies, S. J., Díaz, S., Engelbrecht, B. M. J., Harms, K. E., Hubbell, S. P., Marks, C. O., Ruiz-Jaen, M. C., Salvador, C. M., and Zanne, A. E. (2010). Functional traits and the growth–mortality trade-off in tropical trees. *Ecology*, 91(12):3664–3674.