

Drift dynamics in microbial communities and the effective community size

William T. Sloan ^{1*}, Chioma F. Nnaji,¹ Mary Lunn,²
Thomas P. Curtis ³, Sean D. Colloms ⁴,
Jillian M. Couto ¹, Ameet J. Pinto ⁵,
Stephanie Connelly ¹ and Susan J. Rosser ⁶

¹School of Engineering, University of Glasgow,
University Avenue, Glasgow, G12 8QQ, UK.

²Department of Statistics, University of Oxford, 24-29 St
Giles, Oxford, OX1 3LB, UK.

³School of Engineering, Newcastle University,
Newcastle upon Tyne, NE1 7RU, UK.

⁴Institute of Molecular, Cell and Systems Biology,
University of Glasgow, University Avenue, Glasgow, G12
8QQ, UK.

⁵Department of Civil and Environmental Engineering,
Northeastern University, 360 Huntington Avenue,
Boston, MA, 02115.

⁶School of Biological Sciences, Roger Land Building,
Alexander Crum Brown Road, The King's Buildings,
Edinburgh, EH9 3FF, UK.

Summary

The structure and diversity of all open microbial communities are shaped by individual births, deaths, speciation and immigration events; the precise timings of these events are unknowable and unpredictable. This randomness is manifest as ecological drift in the population dynamics, the importance of which has been a source of debate for decades. There are theoretical reasons to suppose that drift would be imperceptible in large microbial communities, but this is at odds with circumstantial evidence that effects can be seen even in huge, complex communities. To resolve this dichotomy we need to observe dynamics in simple systems where key parameters, like migration, birth and death rates can be directly measured. We monitored the dynamics in the abundance of two genetically modified strains of *Escherichia coli*, with tuneable growth characteristics, that were mixed and continually fed into

10 identical chemostats. We demonstrated that the effects of demographic (non-environmental) stochasticity are very apparent in the dynamics. However, they do not conform to the most parsimonious and commonly applied mathematical models, where each stochastic event is independent. For these simple models to reproduce the observed dynamics we need to invoke an ‘effective community size’, which is smaller than the census community size.

Introduction

It is widely accepted that ecological drift plays a role in shaping all biological communities (Vellend and Agrawal, 2010). However, its relative importance remains a source of debate with potential important consequences, especially for quantifying the uncertainty in strategies to manipulate microbial communities for engineering, medical or agricultural practices (Zhou and Ning, 2017). The doubt about the importance of demographic randomness, which causes drift, is particularly pertinent to microbial communities because, theoretically (Nee, 2005; Ricklefs, 2006), the importance diminishes with increasing population size. Thus, for example, Nee (2005) and Ricklefs (2006) considered the dynamics resulting from a stochastic birth–death model, where each event is independent and randomly distributed in space and time. They showed that in a large population the drift associated with randomness would occur over unfeasibly long timescales and, therefore, would be unobservable and entirely swamped by the environmental effects on the population. Most microbial communities of consequence to us, such as those in soil, the gut, wastewater or infections, comprise huge numbers of individual organisms.

Yet this theoretical argument is at odds with many results from researchers who have just gone ahead and deployed stochastic birth–death processes to describe community assembly (Sloan *et al.*, 2006; Woodcock *et al.*, 2007; Morris *et al.*, 2013; Adair *et al.*, 2018; Ling *et al.*, 2018; Mo *et al.*, 2018). The interest in random effects was initially brought into focus, and simultaneously polarized (McGill, 2003), by neutral theory

Received 8 October, 2020; revised 26 January, 2021; accepted 3 March, 2021. *For correspondence. E-mail william.sloan@glasgow.ac.uk; Tel. +44 (0)141 330 4076.

(Hubbell, 2001), which is an extreme case where only drift is at play in local communities. So, for example, most theoretical manifestations of the purely neutral assumptions yield the results that the abundance of an individual taxa in an isolated community is beta distributed (Sloan *et al.*, 2006) and that the taxa-abundance distribution of all taxa in that community adheres to a Dirichlet distribution (Sloan *et al.*, 2007). It has been shown, using a suite of different methods, that these distributions fit data from a single community or multiple, similar, but geographically distinct communities for single snap-shots in time (Volkov *et al.*, 2003; Sloan *et al.*, 2006; Woodcock *et al.*, 2007). However, fitting stationary taxa abundance distribution to observations will only ever provide circumstantial evidence of the importance of ecological drift. The models may fit but for the wrong reasons; a range of different processes could yield very similar stationary distributions. There are at least as many studies where the neutral models do not fit the data (Clark and McLachlan, 2003; Dornelas *et al.*, 2006). In addition, the assumption of neutrality, where all taxa have the same specific birth and death rate, is at odds with biologists' experience in the laboratory. The fact that neutral models deliver an extreme manifestation of the effects of drift is sometimes overlooked and where neutral models fail to describe patterns in ecology the whole concept of demographic stochasticity and thus drift is dismissed in favour of alternative models built on entirely different axioms (Tilman, 2004). Yet to discard the building block of a random birth–death process as a collateral effect of dismissing the neutral assumption would seem hasty. The drift effects of demographic stochasticity may still be manifest in the dynamics of communities where growth rates are differentiated by species.

The dichotomy of theoretical slow drift in large communities alongside circumstantial experimental evidence of its importance in shaping community composition can only be resolved with richer datasets. In particular, time series of the population dynamics in systems where some of the key parameters that determine the magnitude and speed of the random effects can be directly measured. Time series of abundance of bacteria in communities are quite rare. Ofițeru *et al.* (2010) were able to calibrate a stochastic birth–death process for microbes in a wastewater treatment plant and provide stronger evidence that somehow the demographic stochasticity was indeed important, even in large communities. However, the complexity of that environment made it difficult to discover why the phenomena might be observed in such a large population; the key parameters could only be inferred. Cira *et al.* (2018) created time series using relatively complex synthetic communities of bacteria and again inferred parameter values. In a forest, where a comprehensive time series of tree species exist, again

using model inference alone it appeared that the role of demographic stochasticity was difficult to disentangle from environmental variance (Chisholm *et al.*, 2014). Here, we observe the dynamics in a set of very simple experiments where the environment, as far as is possible, stays constant and the factors that are known to affect the rate of drift are controlled and measured.

Birth, death and immigration rates can be controlled in chemostats, so we describe experiments where we operate 10 identical chemostats in parallel and record the dynamics of two strains of *E. coli*. The strains only differ in having a single copy of a different antibiotic resistance gene added to exactly the same position in their genomes. They behave neutrally when mixed without antibiotics or can be disadvantaged by adding one or other of the antibiotics. We then fit a modest adaptation of the birth–death–immigration model in Ofițeru *et al.* (2010) to the experimentally observed population dynamics to ascertain whether drift explains the dynamics in the chemostats. There are two random effects in our experimental data: the drift, which is a fundamental component of the real dynamics, and experimental errors. Drift is manifest as a random walk where steps are a function of the strain's abundance and successive data in a time series are correlated. Experimental error is manifest as the addition of independent realizations of a normally distributed random variable to the true population numbers at each time point. Our model and fitting method distinguishes between these. We can independently verify the estimate of measurement error from repeat measurements and thus it is possible to reliably quantify the drift. We believe our approach is a uniquely rigorous exploration of drift that is particularly suited to the tractable microbial community dynamics in our experimental setup.

Methods

Two bacterial strains

Truly neutral species would have identical growth rates in all environmental conditions. Unfortunately, it would be extremely difficult to guarantee such equivalence in two naturally occurring species. Thus, we created two genetically modified strains of *E. coli* from a common laboratory parent strain (K-12) by replacing the single-copy non-essential *pepA* gene in its genome with two different antibiotic resistance genes: one conveying chloramphenicol resistance and the other kanamycin resistance. This was achieved through recombination-mediated genetic engineering using bacteriophage lambda-encoded recombination enzymes (Supplementary Note 1), which is reported to be precise and efficient (Datsenko and

Wanner, 2000). Placing the antibiotic resistance gene in the chromosome, rather than on a plasmid, offers more stability. We tested that the single copy *pepA* gene was replaced by only one copy of the antibiotic resistance genes (Supplementary Note 2) and, hence, that we could quantify the abundance of the two modified strains, when they were mixed in known ratios, using qPCR (Supplementary Note 3). This meant we could be confident in qPCR estimates of the abundance of the two populations when mixed and fluctuating in the subsequent chemostat experiments.

In benign conditions, the growth characteristics of the two modified strains should be identical. We showed that this was indeed the case by estimating the growth in triplicate batches for different substrate (tryptone broth) concentrations at 37°C (Supplementary Note 4). The relationship between growth rate and substrate concentration was then characterized using a Monod kinetic model; and the maximum growth rate, $\mu_{\max} \cong 0.85 \text{ h}^{-1}$, and the half-saturation constant, $K_{\text{sat}} \cong 2.5 \text{ g L}^{-1}$, for both strains. We similarly fitted Monod curves to the relationship between substrate concentration and growth rate when the strains were inhibited by the addition of antibiotics. In the presence of $0.5 \mu\text{g ml}^{-1}$ chloramphenicol the growth rate of the chloramphenicol resistant, μ_c , strain is two times that of the kanamycin-resistant strain, μ_k , across a wide range of substrate concentrations (Supplementary Note 4). The kinetic curves also allowed us to design the chemostat setup using a conventional

deterministic mathematical model of chemostats to ensure that the biomass was turned over and retained in the system.

A schematic of the experimental system is given in Fig. 1. It was designed so that different experimental setups could be implemented by simply switching on or off pumps. Three separate experiments were conducted:

1. The first experiment aimed to maximize the chance of observing neutral dynamics, if they occur. In this first experiment, neither of the strains were disadvantaged by the addition of antibiotic, so pump B was not active. The A pumps were active so that the two large chemostats (2 L) were fed with 10 g L^{-1} tryptone-broth at a rate of 0.6 ml min^{-1} which maintained the pure cell cultures at similar densities. These pure cultures were pumped at a rate of 0.55 ml min^{-1} into two separate manifolds and divided evenly between 10 smaller chemostats where they were mixed and continuously stirred. No additional substrate was added and so the mixed cultures in smaller chemostats grew in the residual substrate entering with the bacteria from the large chemostats. The small chemostats overflowed when the volume reached 250 ml and so the volume was held constant at that value and the volumetric flow rate of influent and effluent were similar, not accounting for evaporation. The mass flux of residual substrate ensured that there was growth in the small chemostats so the cell density in the effluent always exceeded that of the influent. The bacteria entering

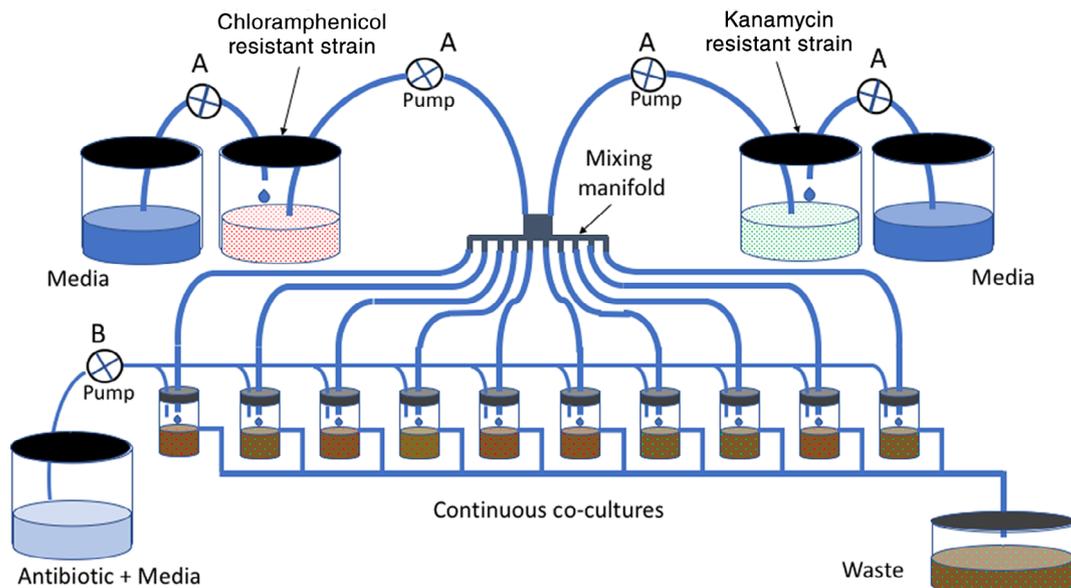


Fig 1. Schematic of the experimental setup. In experiment 1, which potentially creates neutral dynamics in the 10 small co-culture chemostats, substrate and migrants entered via the pumps labelled A; the one labelled B was off. In experiment 2, the initially inoculated chemostats were fed substrate with $0.5 \mu\text{g L}^{-1}$ chloramphenicol via pump B, but no new migrants were allowed to enter, so the A pumps were off. In experiment 3, substrate and migrants entered via the A pumps and $0.5 \mu\text{g L}^{-1}$ chloramphenicol in substrate was fed in via pump B. [Color figure can be viewed at wileyonlinelibrary.com]

the small chemostats will be referred to as immigrants and those in the effluent as emigrants. The entire experimental setup was housed in a temperature-controlled chamber that was maintained at 37°C. The chemostats were operated identically for 7 days. After allowing 2 days for the chemostats to achieve steady cell densities, the large pure-culture chemostats and the effluent from the 10 smaller chemostats were sampled every 2 h during the 5 remaining days (08.00–20.00 h) and the cell density (cells ml⁻¹) was enumerated by QPCR.

- In the second experiment, the 10 smaller communities were inoculated with a 50:50 mix of the two bacterial strains; the stabilizing force of immigration was removed (pumps A switched off); the small chemostats were fed bacteria-free Tryptone-broth substrate (pumps B switched on) with 0.5 µg ml⁻¹ chloramphenicol via a manifold at a rate of 1.1 ml min⁻¹ and so the kanamycin-resistant strain was disadvantaged. Thus, we would expect the chloramphenicol-resistant strain to become monodominant with the dynamics following variable trajectories if demographic stochasticity plays a role and with identical trajectories if it does not. In this experiment, the chemostats were maintained at a volume of 200 ml.
- In the third, and final experiment, the setup was the same experiment as 1, but with the addition of 0.5 µg ml⁻¹ chloramphenicol in tryptone media via pump B at a rate of 0.33 ml min⁻¹ directly into the 10 small chemostats. Also, the small chemostats were kept at a volume of 200 ml. Here, the chloramphenicol-resistant strain should dominate, but the kanamycin-resistant strain ought to be maintained at low abundance by new migrants. It should be possible to ascertain whether any dynamics can be attributed to those induced by the randomness of birth, death immigration events.

The mathematical model is a modest adaptation of that in Ofițeru *et al.* (2010), which derives from Sloan *et al.* (2006) and ultimately Hubbell (2001). It describes a birth–death–immigration process for the co-culture of the two species in small chemostats that are maintained with a constant number of individuals, N_T . For the assemblage to change an individual is selected at random to die or leave the system. The dead individual is replaced by an immigrant from a source community with the probability m or by reproduction by a member of the local community with probability $(1 - m)$. If we take either one of our two strains comprising initially of N individuals, then we can write the probabilities that after one replacement event, there has been an increase by one, decrease by one, or no change (Sloan *et al.*, 2006),

$$\Pr(N+1/N) = \left(\frac{N_T-N}{N_T}\right) \left[mp + (1+\alpha)(1-m)\left(\frac{N}{N_T-1}\right)\right] = b_N, \quad (1)$$

$$\Pr(N-1/N) = \frac{N}{N_T} \left[m(1-p) + (1-\alpha)(1-m)\left(\frac{N_T-N}{N_T-1}\right)\right] = d_N, \quad (2)$$

$$\Pr(N/N) = 1 - b_N - d_N, \quad (3)$$

respectively, where

$$p = \frac{K}{K + K_{\text{other}}}. \quad (4)$$

and K and K_{other} are the number of individuals in the large monoculture of the strain of interest and the other strain respectively. Thus p is the relative abundance of the strain in the influent to the smaller chemostats. α is a parameter that conveys an advantage on the strain by increasing the probability of its birth relative to other strain within the chemostat; when this is set to zero the dynamics are neutral. Note $\Pr(N/N)$ is independent of α . These transition probabilities describe a discrete Markov chain process that can simulate the dynamics of the abundance of the species. In the discrete Markov transition probabilities, each transition represents one replacement event. Simulating the dynamics using this discrete model rapidly becomes cumbersome as N_T gets large. For microbial communities, where N_T is very large indeed, it is completely impractical. So, it is convenient to recast the model as a set of continuous stochastic differential equations (SDEs). The model presented here is very similar to that of Ofițeru *et al.* (2010) but differs slightly. For the complex wastewater treatment plant, where they applied their model, the average time between birth/death random replacements was experimentally intractable and therefore was estimated by calibration. With the careful experiments conducted here, we do know how many replacements occur within a given time period so we can estimate the average time between replacement events. We can also estimate the total population size, N_T , the source community relative abundance, p , and the immigration probability, m , directly from the experimental data. This allows us to give a richer interpretation of the parameters that are delivered by calibrating our SDE. Hence, we relate discrete events in time to continuous real-time in a marginally more intuitive way. We let η be the average time between individual replacement events in the populations and define a scaled time $\tau = t/\eta$. Then in a time period $\Delta\tau = 1$ we can expect one replacement in the community as a whole so that $\Delta\tau$ is, on average, the transition time in the discrete Markov model. In

general, N_T is large we can use the continuous variable $X = \frac{N}{N_T}$ represent the Markov chain process by the SDE (2),

$$dX = M d\tau + \sqrt{V} dW_\tau, \tag{5}$$

where W_τ is a standard Wiener process and for small $\Delta\tau$,

$$E(\Delta X / X(\tau) = x) = M \Delta\tau, \tag{6}$$

and

$$E((\Delta X)^2 / X(\tau) = x) = V \Delta\tau. \tag{7}$$

For the discrete time model, the smallest change that can occur is the replacement of one individual, which we expect to occur in a time $\Delta\tau = 1$. So, during this time period the possible changes in X are $\Delta X = \pm \frac{1}{N_T}$. Consequently, for $x = \frac{N}{N_T}$, we have,

$$E(\Delta X / X(\tau) = x) = \frac{1}{N_T} (b_N - d_N) \tag{8}$$

$$= \frac{1}{N_T} (m(p-x) + 2\alpha(1-m)x(1-x))$$

and

$$E((\Delta X)^2 | X(\tau) = x) = \frac{1}{N_T^2} (2(x(1-x)) + (m(1-2x)(p-x))) \tag{9}$$

which gives us M and V . In Sloan *et al.* (2006), it is argued that the second term in V is an order of magnitude smaller the first and it is reasonable to ignore it. So that Eq. 5 becomes

$$dX = \frac{1}{N_T} (m(p-X) + 2\alpha(1-m)X(1-X)) d\tau + \sqrt{\frac{1}{N_T^2} (2X(1-X))} dW_\tau. \tag{10}$$

Now we can convert to real time using $d\tau = dt/\eta$ so that,

$$dX = \frac{1}{\eta N_T} (m(p-X) + 2\alpha(1-m)X(1-X)) dt + \frac{1}{\sqrt{\eta N_T^2}} \sqrt{(2X(1-X))} dW_t. \tag{11}$$

or

$$dX = f(X, p; \theta_1, \theta_2) + g(X; \theta_3) dW_t, \tag{12}$$

where

$$f(X, p; \theta_1, \theta_2) = \theta_1(p-X) + \theta_2 X(1-X),$$

$$g(X; \theta_3) = \theta_3 \sqrt{(2X(1-X))},$$

$\theta_1 = \frac{m dt}{\eta N_T}$, $\theta_2 = \frac{2\alpha(1-m) dt}{\eta N_T}$ and $\theta_3 = \frac{1}{\sqrt{\eta N_T^2}}$. This then constitutes our SDE that we hope will describe the dynamics in relative abundance of either one of our bacterial strains.

The fitting method

To calibrate the model we have n observations of the relative abundance $\{X_i, i = 1, n\}$ that were made every 2 h, $dt = 2$, in all of the small chemostats, from 08:00 to 20:00 h daily. We also have observations of the relative abundance $\{p_i, i = 1, n\}$ in the large pure cultures and we can calculate $dX_i = X_{i+1} - X_i$ for all observations except the couplets that span the night, where the time gap was 12 h, which is too-long. The relative abundances were quantified using qPCR with the same standard curve which should minimize the variability in measurements; nonetheless, there will be some measurement error. So, for our measurements,

$$dX_i = f(X_i, p_i; \theta_1, \theta_2) + g(X_i; \theta_3) dW_{t,i} + \varepsilon_i \tag{13}$$

where $dW_{t,i}$ is the i th realization of the standard Weiner process $dW_{t,i} \sim \sqrt{dt} N(0, 1)$ at ε_i is a realization of the measurement error in X . We make the assumption that measurement noise, ε_i , has the effect of adding a constant, σ_ε , to standard deviation in the process. To fit the model we use maximum likelihood estimation (MLE), where the log-likelihood function for the parameters given the data is

$$\mathcal{L}(\theta_1, \theta_2, \theta_3, \sigma_\varepsilon; \{X_i, dX_i, p_i; i = 1, n\})$$

$$= - \left(\sum_{i=1}^n \left(\frac{(dX_i - f(X_i, P_i; \theta_1, \theta_2) dt)^2}{2 (dtg(X_i; \theta_3)^2 + \sigma_\varepsilon^2)} \right) + \ln \left(\sqrt{2\pi (dtg(X_i; \theta_3)^2 + \sigma_\varepsilon^2)} \right) \right) \tag{14}$$

This allows us to make an estimate of all the parameters including the measurement noise σ_ε based solely on the time-series data, which would not be possible using, for example, least squares fitting. Then, merely to gain more information on the quality of the model fit, the MLE estimate of σ_ε is fixed in Eq. 13, which allows us to use ordinary least squares, as in Ofiteru *et al.* (2010), to deliver goodness-of-fit statistics.

Results

The time series of relative abundance for both bacterial strains in all experiments and reactors are given in supplementary materials. We have highlighted the generic behaviour of the dynamics in relative abundance in the experiments by plotted graphs for an arbitrarily selected subset of the reactors. Plotting all reactors for all time points would lead to too many, or overly complicated, figures and plotting the averages across all reactors would defeat the purpose of highlighting the variability between reactors.

For experiment 1, the time series of bacterial abundance in the pure culture of the chloramphenicol-resistant strain in the large monoculture chemostat and for two of the smaller co-culture chemostats are presented in Fig. 1A. We have not displayed the relative abundance of the kanamycin strain because it is merely one minus the relative abundance of the chloramphenicol strain. The graph shows that for both of these small chemostats the relative abundance fluctuates around the relative abundance of the large monoculture from which immigrants are drawn. The dynamics in the fluctuation about the source community abundance seems to be independent in the two reactors. This observation is similar for all 10 of the small co-culture chemostats; Fig. 1B demonstrates this by plotting the dynamics in abundance for 1 day for all 10 chemostats. So, whilst the abundance of the source of immigrants strongly influences the abundance in the local co-culture chemostats, there are serially correlated dynamics occurring independently in each of the vessels, despite the environmental conditions being the same in each. This is redolent of ecological drift.

In experiment 2 we inoculate the chemostats and cut-off immigration, so that the probability of an immigration event $m = 0$. The addition of $0.5 \mu\text{g L}^{-1}$ of chloramphenicol to each of the co-culture chemostats should convey an advantage to the chloramphenicol-resistant strain. In Fig. 1C we see that the experimental results, presented here for just three arbitrarily selected chemostats, bear this out. The chloramphenicol strain ultimately becomes dominant in all chemostats, but the route towards dominance differs in the reactors. Indeed, in two of the reactors displayed in Fig. 1C the chloramphenicol-resistant strain initially fares less well than the kanamycin-resistant strain. Again, the dynamics vary from reactor to reactor despite the environmental conditions being identical.

In experiment 3, we reintroduce migration from the large monoculture chemostats and see, for two arbitrarily selected reactors (Fig. 2D), that whilst the chloramphenicol strain dominates it does not become monodominant; the kanamycin-resistant strain is retained in the chemostat by the steady stream of migrants, despite being disadvantaged by the antibiotic.

Our use of qPCR allowed us to estimate the total population size, N_T . Organisms enter the reactor and grow, so the density in the reactor is higher than in the influent, however, at a steady state that density stays constant in time. Therefore, the number of replacements either by births or by immigration in a unit of time must be equal to the number leaving in the effluent. So, if the density of bacteria in the effluent is N_{eff} and the volumetric flow rate is Q then the total number of replacements in the community per unit time is simply QN_{eff} and thus the time between replacement is simply $\eta = \frac{1}{QN_{\text{eff}}}$. If N_{in} is the density of immigrants in the influent, then the number of immigrants per unit time is QN_{in} and so the probability that when a replacement event occurs it is by an immigrant is just $m = \frac{QN_{\text{in}}}{QN_{\text{eff}}} = \frac{N_{\text{in}}}{N_{\text{eff}}}$. In supplementary materials, we see that in the presence of $0.5 \mu\text{g L}^{-1}$ chloramphenicol the growth rate of the chloramphenicol resistant, μ_c , strain is two times that of the kanamycin-resistant strain, μ_k , across a wide range of substrate concentrations. Thus, $\frac{\mu_c}{\mu_k} = \frac{1+\alpha}{1-\alpha} = 2$ which means that $\alpha = 0.33$. So, estimates of m , η , N_T and α can be calculated directly from the experiments and are given in Table 1.

We can also approximate the standard deviation of the measurement error, σ_e . The abundances of each of the antibiotic-resistant genes were enumerated at all time points in the small reactors using qPCR with three technical replicates. All qPCR estimates for each gene used the same standard curves. With three estimates of the kanamycin-resistant gene and three of the chloramphenicol then we have nine possible combinations to estimate the total abundance of bacteria in the samples at each time point. And then if we consider all possible ratios of, for example, the three measurements of the abundance of chloramphenicol strain with the nine estimates of total abundance then there 27 possible values. At each time point, we calculated the mean and standard deviation of the possible values and then we calculated the mean of the standard deviation across all time points to obtain a crude estimate of σ_e . We obtained a value of $\sigma_e \simeq 0.07$ by averaging across all the data (Table 1).

One set of parameters was calibrated for each experiment using the time series from all the chemostats (Table 2). The p -values for the least squares regression fit demonstrate that the model parameters are significantly different from the null hypothesis of them being zero (no trend). Bearing in mind that we are fitting a SDE with the aim of explaining the variability between reactors we would not necessarily expect a high R -square value for the fit, which is the case, especially for experiment 2 where there is no stabilizing influence of immigration. Overall, though, the significances of the over-all model fits are high. It is more important to test the null hypothesis that the weighted residuals are normally distributed.

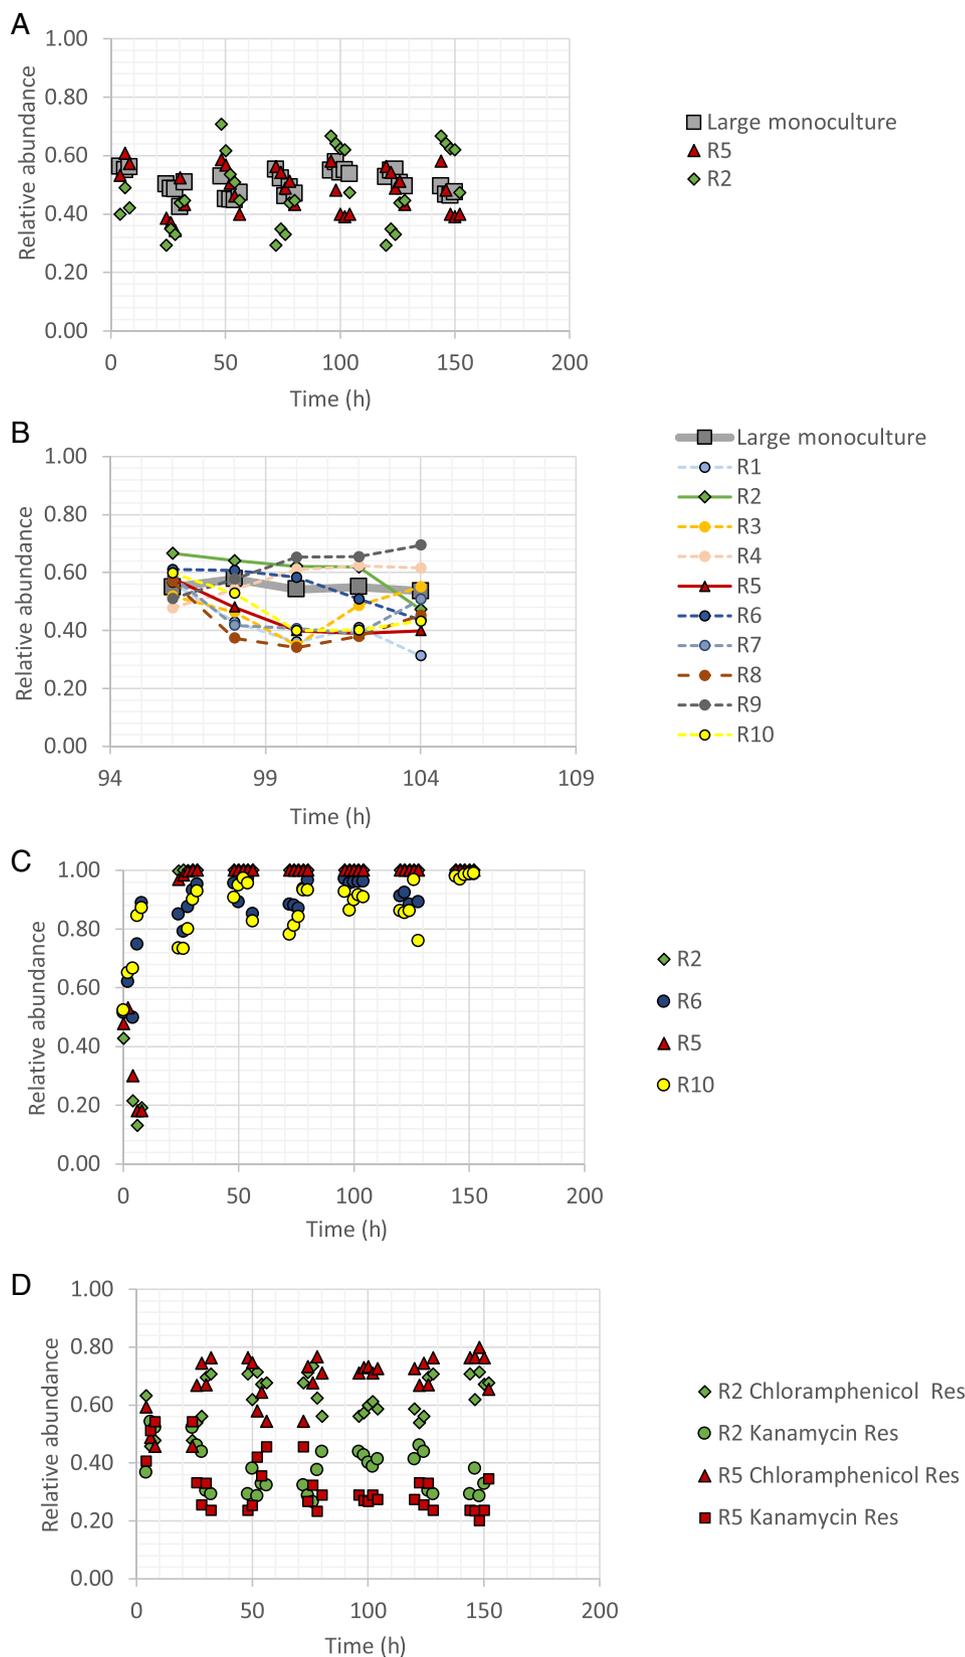


Fig 2. A. Experiment 1, where we might expect neutral dynamics. A graph of the dynamics in relative abundance of the chloramphenicol-resistant strain in the large monoculture that acts as a source of immigrants and in two arbitrarily selected smaller co-culture chemostats (2 & 5). B. The relative abundance of the chloramphenicol-resistant strain in all 10 of the co-culture chemostats for a 10 h period during experiment 1. C. Experiment 2, where immigration to the small chemostats is cut off and chloramphenicol is added, which gives the chloramphenicol-resistant strain an advantage. The plot shows the chloramphenicol-resistant strain taking a variety of trajectories towards becoming monodominant in an arbitrary selection of four reactors (2, 5, 6 & 10). D. Experiment 3, where immigrants are equally likely to be of either strain, but the addition of chloramphenicol conveys an advantage on the chloramphenicol-resistant strain. A graph of both the relative abundance of chloramphenicol and kanamycin-resistant strains in two of the co-culture chemostats. The chloramphenicol-resistant strain is consistently more abundant but the dynamics vary between reactors. [Color figure can be viewed at wileyonlinelibrary.com]

For the case with no immigration, where the chloramphenicol-resistant strain becomes monodominant, then 1 and 0 are absorbing boundaries where the residuals become very close to zero; which erroneously skews the distribution of residuals. So we test the fit for the region just before the process first hits the boundary, which is analogous to first passage time problem in survival analysis (Lee and Whitmore, 2006), and thus we only consider residuals where the abundance lies between 0.05 and 0.95. Using the Anderson–Darling test, we fail to reject the null hypothesis of normality at the 95% level. So, it would appear that the SDEs are doing a good job of capturing the noisy dynamics across all the reactors.

Discussion

There was a high degree of variability in the dynamics of the abundance of the two strains across identically operated chemostats. We were able to fit the parameters of our SDEs using the time series data from all three experiments with a high degree of significance. This would suggest that we are observing neutral dynamics with immigration in experiment 1, selection and drift in experiment 2 and selection and drift mediated by immigration in experiment 3. Precisely what we anticipated. So, it would appear, on this evidence, that demographic stochasticity does play an important role in the dynamics, whether or not the species have identical growth rates.

The beauty of the simple experimental system we have used here is that we can actually measure the key

Table 1. The values of the variables that form the parameters in the birth–death–immigration model, estimated directly from the experimental data.

Experiment	m	η	N_T	α	σ_e
1.	0.68	1.19×10^{-11}	2.26×10^{11}		0.07
2.	–	4.58×10^{-11}	6.61×10^{10}	0.33	0.07
3.	0.31	6.29×10^{-12}	3.7×10^{11}	0.33	0.07

Table 2. The parameter estimates obtained by calibrating the model using the time series of relative abundance for each of the three experiments.

Experiment	θ_1	θ_2	θ_3	σ_e	R^2	p -value model fit .v. no trend	Normally distributed weighted residuals (AD test)
1. Parameter estimate	0.31	–	0.026	0.071	0.17	2.96×10^{-12}	✓
p -value	8.06×10^{-12}	–	–	–	–	–	–
2. Parameter estimate	–	0.16	0.051	0.077	0.07	3.33×10^{-5}	×All residuals
p -value	–	3.32×10^{-5}	–	–	–	–	✓Residuals away from adsorbing boundary
3. Parameter estimate	0.32	0.29	0.041	0.080	0.27	3.59×10^{-18}	✓
p -value	1.20×10^{-18}	6.45×10^{-15}	–	–	–	–	–

variables in the experiment and compare them to our calibrated model parameters. First, the calibrated estimates of the standard deviation in measurement noise are in close agreement with that estimated directly from the triplicate qPCR measurements. For the other parameters, it is important to note that in our mathematical model, and indeed, in every other manifestation of random birth–death population models, the parameters we calibrate are compound combinations of the fundamental parameters that describe the process. So, for example, here $\theta_1 = \frac{m \alpha}{\eta N_T}$. If we estimate the compound parameters for our model using the variables, m , N_T , η and α , measured directly in our experiments (Table 1) do they match the calibrated model parameters (Table 2)? Table 3 makes that comparison.

For parameters θ_1 and θ_2 there is a reasonably good match between the calibrated and experimentally derived estimates. However, values of θ_3 estimated by calibration are several orders of magnitude larger than that are calculated from the direct measured variables. θ_3 determines the size of the random step in the stochastic process and hence the timescales over which stochastic demography becomes apparent.

So, this takes us straight back to the heart of the dilemma that we outlined in the introduction. As Nee (2005) and Ricklefs (2006) suggested, if we use the actual measured population size, immigration probability and replacement rate the model predicts the stochastic contribution to dynamics that is small and would only become apparent over long timescales. Yet, the observed dynamics are highly variable and are well represented by the SDEs albeit with the calibrated parameter that determines the speed of dynamics, θ_3 , much larger than we would expect. So, in our tightly controlled experiments, we see exactly the dichotomy that we described earlier, based on birth–death processes calibrated for more complex environments. This is frustrating. There is a tantalizing alignment between the models and the experimental results that is only undermined by the failure to match one important parameter. Thus, it appears that our underlying model (Eqs. 1–3), which is

Table 3. Comparison between the compound parameters estimated by calibrating the model and those derived from the experimental variables.

Experiment	θ_1	θ_2	θ_3	σ_e
1. Parameter calibrated	0.31	–	2.6×10^{-2}	0.071
Parameter from experimental variables	0.35	–	1.1×10^{-6}	0.070
2. Parameter estimates	–	0.16	5.1×10^{-2}	0.077
Parameter from experimental variables	–	0.21	2.2×10^{-6}	0.070
3. Parameter estimate	0.32	0.29	4.1×10^{-2}	0.080
Parameter from experimental variables	0.29	0.19	1.7×10^{-6}	0.070

the basis of all birth–death processes in ecology is fundamentally flawed.

Before dismissing birth–death processes for our simple microbial communities it is worth considering their application in the related field of population genetics. There the central premise is that mutations in the genomes of a population are inevitable, but the precise timing and location of the mutations defy characterization. So, the mathematics of population genetics, whilst continually advancing, has adhered to the idea that inheritance is an intrinsically stochastic process. The simplest models encapsulate a few fundamental truths for all organisms; they are born, they die, they inherit mutations and they accumulate new ones at random. Selection is manifest as a tweak on the probability that one mutant preferentially reproduces over others. Thus population genetics has simple random birth–death processes at its heart and these are deployed to great effect in range of practical applications including human medicine, animal husbandry, epidemiology and agriculture (Pirchner, 1969). This is the case even although the model predictions, in their rawest form, always deviate from observations. Thus, for example, the models always significantly underestimate the variance in the frequency of alleles in successive populations (Wright, 1931) when applied using the census population number. This does not deter geneticists; confident that the fundamental stochastic processes are at play, they retain the birth–death–mutation model and introduce the idea of an effective population size. That is, a fictitious smaller population size would see the model produce the correct results. This is not as much of a ‘fix’ as it might first seem because there is a wide variety of biological and demographic reasons why it is reasonable to assume that not all of the census population will contribute to the genetic diversity of the next generation. Thus effective population sizes in clonal populations of bacteria estimated using allele frequencies have small effective population sizes (Fraser *et al.*, 2007).

Can we deploy a similar approach here and, based on the assumption that the random birth–death process must be occurring, make sense of the rapid drift by describing an effective parameter that makes biological sense? The model assumes that every organism that is

replaced is selected independently at random from the population. Suppose instead that there was some spatial or temporal correlation in the replacements. This is not unreasonable since, despite our best efforts, it is impossible to perfectly mix a reactor and so when an individual leaves the system it departs at the same time as at least a few of its neighbours. Furthermore, when a bacterium divides the offspring are collocated, they are not instantaneously transported to randomly selected locations in the chemostat. In addition, with any bacterial community some degree of flocculation is common.

Our model does not explicitly represent space. However, we can explore what size of spatial or temporal grouping, or ‘clumping’, might yield the stochastic dynamics we observe. If the cells in the chemostat reside in clumps, then they would exit the chemostats in those clumps. So if N_{ec} is the total number of clumps and η_{ec} is the average time between clumps leaving then, $N_{ec}\eta_{ec} = N_T\eta$; the turnover in the population is the same it merely occurs in clumps rather than individuals. Note then that the first two parameters in our model remain unchanged if we replace N_T with N_{ec} and η with η_e no matter what their value. If we assume that $N_{ec}\eta_{ec}$ is given by the experimentally defined $N_T\eta$ then we can now calculate the value of N_{ec} that would yield the calibrated value of θ_3 , which for experiment 1, 2 and 3 would be 598, 126 and 255 respectively; which is tiny in comparison to the number of individuals. Now consider the number of organisms that would constitute a clump, N_T/N_{ec} , which is 5.56×10^8 , 5.21×10^8 and 14.4×10^8 , similar across all three experiments. These seem large, but the density of organisms is very high in the chemostats. Indeed, if the organisms were evenly distributed in the liquid then one of these clumps would only occupy approximately 0.5 ml. If they were correlated in time then it is the number of organisms that would leave the reactors in approximately 30 s. It, therefore, does not seem unreasonable to us that spatial or temporal correlation in the replacement events might be the reason that we seem to observe stochastic dynamics of a magnitude far higher than we would expect if events were truly randomly distributed as they are in our model. We suspect that the spatial correlation of events will be equally important to stochastic dynamics in other environments. Thus,

predators will plough through a population, consuming bacteria and their neighbours rather than picking up individuals at random from across the population. Infiltration of rainfall in soils, attachment to substratum in a reactor or indeed in structured biofilms will all promote correlated random events. Also, it may explain the parameter values that emerge from studies where neutral models are calibrated using stationary taxa abundance distributions. In Etienne and Alonso (2005) description of the stationary taxa-abundance distribution for the neutral case, a 'fundamental dispersal number' is calibrated, and in Sloan *et al.*'s (2006) adaption for microbial communities, a similar compound parameter $N_T m$ is calibrated. Typically the calibrated values of $N_T m$ are in the range of 10s to 10 000s depending on the clade of bacteria (Sloan *et al.*, 2006; Woodcock *et al.*, 2007; Morris *et al.*, 2013; Adair *et al.*, 2018), which in large populations would suggest very low immigration probabilities. But if N_T were replaced by a much smaller effective community size, N_{ec} , then immigration could be orders of magnitude higher and perhaps more realistic for the systems.

The composition and dynamics of any open microbial community are ultimately shaped by four things (Vellend and Agrawal, 2010): diversification, dispersal, selection and drift. If we had a perfect knowledge of every individual organism in a community, from its position to metabolic function and state, if we could anticipate precisely where and when migrants entered or departed and if we had an exact map of the environment now and in the future then we might, in theory, be able to predict the effects of these processes on the functioning and the fate of species and their diversity in microbial communities. However, we do not have this information even for simple communities, let alone the complex communities that we routinely rely on for important biotechnologies in, for example, agriculture and water treatment. Furthermore, it is not clear if we will ever be able to routinely collect such information or choose to deploy the computing power to process it. So if we are to use models to predict critical elements of the behaviour of important communities then must we accept models that are gross-simplifications of the real world but, importantly, capture stable aspects of the system (Cox, 1995). In population genetics, this perspective is deeply ingrained. Thus, tractable models that describe the drift and selection of alleles are built on often overly simplified but convenient assumptions like, random mating, simultaneous birth of each new generation, constant population size and equal numbers of children per parent. When applied to real-world populations, in order for the 'ideal' model predictions to match data it is usually necessary to use an effective population size that is smaller than the census population size. The effective population size has emerged as one of the most important characteristics in predicting within-species allele dynamics and diversification.

In microbial ecology, we have no such convention. We are less inclined to accept such abstractions and more inclined to limit ourselves to first-order descriptors that take no account of any underlying conceptual model, even though we have some consensus (Vellend and Agrawal, 2010). Here we have demonstrated that by adopting similar effective community sizes, N_{ec} , which are orders of magnitude lower than the census population the dynamics in all three experimental scenarios, neutral, with selection, and totally insular, the dynamics are consistent with an idealized community where all births, replacements and immigration events are independent. Thus, using the effective parameter allows us to make vital predictions on the dynamics of species in the community and the challenge then becomes determining this fundamental characteristic for different consortia and environments. N_{ec} , we believe, can be deployed as a fundamental measure to help characterize the dynamics of mixed bacterial consortia in all sorts of important environments from wastewater reactors to the human gut or soils and ultimately be as useful as effective population size has been in genetics.

Acknowledgements

This research was funded by two UK, Engineering and Physical Sciences Research Council grants, EP/H009604/1 and EP/K038885/1, and by a Royal Academy of Engineering Chair in Emerging Technology for WTS.

References

- Adair, K.L., Wilson, M., Bost, A., and Douglas, A.E. (2018) Microbial community assembly in wild populations of the fruit fly *Drosophila melanogaster*. *ISME J* **12**: 959–972.
- Chisholm, R.A., Condit, R., Rahman, K.A., Baker, P.J., Bunyavejchewin, S., Chen, Y.Y., *et al.* (2014) Temporal variability of forest communities: empirical estimates of population change in 4000 tree species. *Ecol Lett* **17**: 855–865.
- Clark, J.S., and McLachlan, J.S. (2003) Stability of forest biodiversity. *Nature* **423**: 635–638.
- Cox, D. (1995) Comment on "Model uncertainty, data mining and statistical inference". *J R Stat Soc*, **158**, 455–456.
- Cira, N.J., Pearce, M.T., and Quake, S.R. (2018). Neutral and selective dynamics in a synthetic microbial community. *Proceedings of the National Academy of Sciences*, **115**(42): E9842–E9848.
- Datsenko, K.A., and Wanner, B.L. (2000) One-step inactivation of chromosomal genes in *Escherichia coli* K-12 using PCR products. *Proc Natl Acad Sci U S A* **97**: 6640–6645.
- Dornelas, M., Connolly, S.R., and Hughes, T.P. (2006) Coral reef diversity refutes the neutral theory of biodiversity. *Nature* **440**: 80–82.
- Etienne, R.S., and Alonso, D. (2005) A dispersal-limited sampling theory for species and alleles. *Ecol Lett* **8**: 1147–1156.
- Fraser, C., Hanage, W.P., and Spratt, B.G. (2007) Recombination and the nature of bacterial speciation. *Science* **315**: 476–480.

- Hubbell, S.P. (2001) *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton: Princeton University Press.
- Lee, M.-L.T., and Whitmore, G.A. (2006) Threshold regression for survival analysis: modeling event times by a stochastic process reaching a boundary. *Stat Sci* **21**: 501–513.
- Ling, F., Whitaker, R., LeChevallier, M.W., and Liu, W.-T. (2018) Drinking water microbiome assembly induced by water stagnation. *ISME J* **12**: 1520–1531.
- McGill, B.J. (2003) A test of the unified neutral theory of biodiversity. *Nature* **422**: 881–885.
- Mo, Y., Zhang, W., Yang, J., Lin, Y., Yu, Z., and Lin, S. (2018) Biogeographic patterns of abundant and rare bacterioplankton in three subtropical bays resulting from selective and neutral processes. *ISME J* **12**: 2198–2210.
- Morris, A., Beck, J., Schloss, P., Campbell, T., Crothers, K., Curtis, J., et al. (2013) Comparison of the respiratory microbiome in healthy nonsmokers and smokers. *Am J Respir Crit Care Med* **187**: 1067–1075.
- Nee, S. (2005) The neutral theory of biodiversity: do the numbers add up? *Funct Ecol* **19**: 173–176.
- Ofițeru, I.D., Lunn, M., Curtis, T.P., Wells, G.F., Criddle, C.S., Francis, C.A., and Sloan, W.T. (2010) Combined niche and neutral effects in a microbial wastewater treatment community. *Proc Natl Acad Sci U S A* **107**: 15345–15350.
- Pirchner, F. (1969) *Population Genetics in Animal Breeding*. New York: Springer US.
- Ricklefs, R.E. (2006) The unified neutral theory of biodiversity: do the numbers add up? *Ecology* **87**: 1424–1431.
- Sloan, W.T., Lunn, M., Woodcock, S., Head, I.M., Nee, S., and Curtis, T.P. (2006) Quantifying the roles of immigration and chance in shaping prokaryote community structure. *Environ Microbiol* **8**: 732–740.
- Sloan, W.T., Woodcock, S., Lunn, M., Head, I.M., and Curtis, T.P. (2007) Modeling taxa-abundance distributions in microbial communities using environmental sequence data. *Microb Ecol* **53**: 443–455.
- Tilman, D. (2004) Niche tradeoffs, neutrality, and community structure: a stochastic theory of resource competition, invasion, and community assembly. *Proc Natl Acad Sci U S A* **101**: 10854–10861.
- Vellend, M., and Agrawal, A. (2010) Conceptual synthesis in community ecology. *Q Rev Biol* **85**: 183–206.
- Volkov, I., Banavar, J.R., Hubbell, S.P., and Maritan, A. (2003) Neutral theory and relative species abundance in ecology. *Nature* **424**: 1035–1037.
- Woodcock, S., van der Gast, C.J., Bell, T., Lunn, M., Curtis, T.P., Head, I.M., and Sloan, W.T. (2007) Neutral assembly of bacterial communities. *FEMS Microbiol Ecol* **62**: 171–180.
- Wright, S. (1931) Evolution in Mendelian populations. *Genetics* **16**: 97–159.
- Zhou, J., and Ning, D. (2017) Stochastic community assembly: does it matter in microbial ecology? *Microbiol Mol Biol Rev* **81**(4), e00002–e000017.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web-site:

Supplementary S1. Note 1 (SN1).

Supplementary S2. Note 2 (SN2).

Supplementary S3. Note 3 (SN3).

Supplementary S4. Note 4 (SN4).