The population genetics of spatial sorting

Ben L. Phillips[1] and T. Alex Perkins[2]

1. School of BioSciences, University of Melbourne Parkville VIC 3010 Australia. phillipsb@unimelb.edu.au
2. Department of Biological Sciences and Eck Institute for Global Health, University of Notre Dame, Notre Dame, IN USA

19 Feb 2018

Abstract

Recently it has been claimed that natural selection has a "shy younger sibling" called spatial sorting. Spatial sorting is the process whereby variation in dispersal ability is sorted along density clines and will, in nature, often be a transient phenomenon. Spatial sorting is, however, persistent on invasion fronts, where its effect cannot be ignored, causing rapid evolution of traits related to dispersal. While spatial sorting is described in several elegant models, these models require a high level of mathematical sophistication and are not accessible to most evolutionary biologists or their students. Here we incorporate spatial sorting into the classic haploid and diploid models of natural selection. We show that spatial sorting can be conceptualized precisely as selection operating through space rather than (as with natural selection) time, and that genotypes can be viewed as having both spatial and temporal aspects of fitness. The resultant model is strikingly similar to the classic models of natural selection. This similarity renders the model easy to understand (and to teach), but also suggests that many established theoretical results around natural selection will apply equally to spatial sorting.

Introduction

The modern evolutionary synthesis formalized evolutionary theory, and reconciled particulate inheritance with Darwin's big idea: natural selection. A vast edifice of theory has now been built, and useful entry points into this theory include the basic haploid and diploid models of natural selection. These models, originally conceived free of population dynamics and expressed as either a change in allele ratios or allele frequencies over time (Haldane 1924; Wright 1931), have become the standard models by which students are introduced to the theory of natural selection (e.g., Hartl et al. 1997) and over the years have become more explicitly linked to population dynamics (e.g., Crow and Kimura 1970). A particularly thorough and lucid derivation is given in Otto and Day (2007), in which we see the frequency of the A allele, p, following births and deaths given by

$$p(t+1) = \frac{W_A p(t)}{W_A p(t) + W_a(1 - p(t))} \tag{1}$$

for a haploid system, and

$$p(t+1) = \frac{W_{AA} p^2(t) + W_{Aa} p(t)(1 - p(t))}{W_{AA} p^2(t) + 2W_{Aa} p(t)(1 - p(t)) + W_{aa}(1 - p(t))^2} \tag{2}$$

for a diploid system. In these models, the W terms represent the fitnesses (per capita balance of births and deaths) of individuals carrying each genotype, and the denominator in each case is the mean fitness of the population.

These models capture the basic idea of natural selection and have been used as a launching pad for many theoretical forays in population genetics. Recently, however, it has been pointed out that natural selection may not be the only mechanism for directed evolutionary change, and that natural selection may have a "shy younger sibling" – spatial sorting – a spatial analogue of natural selection (Shine et al. 2011). Spatial sorting arises when there is a cline in density and variation for dispersal ability. When low-density habitat is colonized, it tends to be colonized disproportionately by individuals with superior dispersal abilities, simply because they tend to arrive sooner. Individuals are sorted through space according to dispersal ability, resulting in a directed evolutionary shift in dispersal through space and time. The idea was first adumbrated by Cwynar and MacDonald (1987) and has been re-discovered and refined by various authors since (e.g., Olivieri et al. 1995; Travis and Dytham 2002; Hughes et al. 2003; Phillips et al. 2008; Burton et al. 2010), eventually being dubbed "spatial sorting" in 2011 (Shine et al. 2011).

In many contexts, spatial sorting is likely so fleetingly transient as to be easily missed. Despite this transience, its aggregate effect can still be observed on the evolution of dispersal rates in metapopulations (Van Valen

1971; Olivieri et al. 1995). In the context of spatially expanding populations, however, spatial sorting is persistent on the invasion front for the duration of spread and for some time after. In this situation, the invasion front colonizes unoccupied space every generation, spatial sorting is sustained over time on the invasion front, and its effects are difficult to ignore. The clearest natural example of this comes from the spread of cane toads across northern Australia. Here, evolved shifts in dispersal ability contributed to a five-fold increase in invasion speed over 70 generations (Phillips et al. 2008, 2010; Perkins et al. 2013). Numerous other natural examples have come to light, ranging across taxa from insects to plants (e.g., Cwynar and MacDonald 1987; Simmons and Thomas 2004; Lombaert et al. 2014). More compelling still, are a growing list of laboratory studies showing repeatable evolutionary shifts in dispersal on invasion fronts (e.g., Ditmarsch et al. 2013; Fronhofer and Altermatt 2015). Two recent laboratory studies on beetles also unequivocally demonstrate that these evolutionary shifts are due to spatial sorting (Ochocki and Miller 2017; Weiss-Lehman et al. 2017).

The original theoretical arguments behind spatial sorting on invasion fronts were illustrated with individual-based simulation models (e.g., Travis and Dytham 2002; Hughes et al. 2003; Phillips et al. 2008; Burton et al. 2010). While these painted a convincing picture and pointed to an interesting range of theoretical possibilities (including the possibility that spatial sorting could cause a reduction in fitness), those findings may have come at the cost of generality. How are we to know that these results were not sensitive to the number of individuals simulated or the width of spatial corridors used in these simulations? The less sensitive to these choices the results from these models are, the more readily spatial sorting can be assimilated into evolutionary theory.

More recently, theoreticians have worked to integrate variation in dispersal into equation-based models of biological spread. Although mathematically challenging, these equation-based models have the potential to generate cleaner notions of how certain biological factors modulate the dynamics of spatial sorting. These models treat dispersal as a quantitative trait embedded within an integral projection (e.g., Ellner and Schreiber 2012), integrodifference (e.g., Perkins et al. 2013), or partial differential (e.g., Perkins et al. 2016) model. Additionally, an important generalisation of Fisher's (Fisher 1937) reaction-diffusion model has been conceived that elegantly encapsulates the process of spatial sorting by allowing the model's "diffusion coefficient" to evolve (Benichou et al. 2012; Bouin et al. 2012; Bouin and Calvez 2014). This new class of reaction-diffusion model is currently the focus of intense and productive theoretical work, much of which necessarily involves rather advanced mathematics and is beyond the grasp of many biologists.

In this paper, we introduce a simple conceptual arrangement that allows us to introduce spatial sorting into foundational models of evolutionary biology: the haploid and diploid selection models. The intention is to demonstrate that spatial sorting can be conceptualized precisely as selection operating through space, rather than time, and that this conceptualization leads to a simple generalisation of the haploid and diploid selection models, which are models that are easy to understand and which most biologists will have encountered during their training.

The conceptual arrangement

Space and time are discrete. We imagine a one-dimensional spatial lattice of large size, with position on the lattice denoted by x (see Fig. 1). At a given time t, birth, death, and dispersal occur, in this order.

At $t = 0$, we imagine that only the left-most patch in our lattice, at $x = 0$, is occupied, containing a large number, n, individuals. All other patches are empty. We assume nearest-neighbor dispersal, which implies that if the dispersal rate is not zero, patch 1 will be occupied at time 1, patch 2 at time 2, and so on. The model tracks the dynamics of this vanguard population as it moves through space, at $x = t$ (Fig. 1). As long as the product of the mean population growth rate and mean dispersal rate is greater than 1, then this frontal deme will increase in size over time.

Figure 1: The spatiotemporal arrangement of the model. Filled black circles show space being occupied as time progresses. The model tracks the dynamics of the foremost patch at $x = t$, shown by the large outer circles and the diagonal arrows.

Haploid dynamics

In a haploid model with discrete generations, we imagine two alleles, A and a. Each allele has both a spatial and temporal aspect to its fitness. W denotes temporal fitness: the per capita number of surviving offspring at $t + 1$, such that $W = (1 + b)(1 - d)$, where b is the expected number of offspring per capita in each time interval and d is the probability of death in each time interval. V denotes spatial fitness, which we define as the probability of dispersal to $x + 1$.

Following reproduction but before dispersal, the number of individuals of each genotype at time t and vanguard patch x is

$$n'_A(t, x) = W_A n_A(t, x) \tag{3}$$
$$n'_a(t, x) = W_a n_a(t, x). \tag{4}$$

Following dispersal, the number of individuals of each genotype at time $t + 1$ and newly occupied patch $x + 1$ is

$$n_A(t + 1, x + 1) = V_A W_A n_A(t, x) \tag{5}$$
$$n_a(t + 1, x + 1) = V_a W_a n_a(t, x). \tag{6}$$

If we let $p(t, x) = \frac{n_A(t,x)}{n_A(t,x) + n_a(t,x)}$, then

$$p(t + 1, x + 1) = \frac{V_A W_A p(t, x)}{V_A W_A p(t, x) + V_a W_a (1 - p(t, x))}, \tag{7}$$

which is similar to the standard haploid model of selection (eqn. 1), except that we have explicitly incorporated *spatiotemporal fitness* via $V_A W_A$ and $V_a W_a$. In the event that temporal aspects of fitness are identical across the two alleles (that is, the two alleles have equal reproductive fitness), spatial fitness drives dynamics; in the event that spatial fitness is identical between the two alleles (that is, the two alleles have equal dispersal tendency), we recover the standard haploid model of natural selection.

Similar to the standard model, eqn. 7 has two equilibria: at $p = 0$ (stable when $\frac{V_A W_A}{V_a W_a} < 1$) and $p = 1$ (stable when $\frac{V_A W_A}{V_a W_a} > 1$). These stability criteria point to the strong interaction between temporal and spatial aspects of fitness, because we can think of them as products of two relative fitnesses, one spatial and one temporal. The implication of this result is that the vanguard will eventually be dominated by one genotype or the other, provided that spatial expansion continues for long enough.

Diploid dynamics

In a diploid model we have three genotypes, AA, Aa, and aa, and a total population size $n(t, x) = n_{AA}(t, x) + n_{Aa}(t, x) + n_{aa}(t, x)$. Following birth, death, and dispersal, the numbers of the three genotypes at time $t + 1$ and patch $x + 1$ follow

$$n_{AA}(t + 1, x + 1) = V_{AA} W_{AA} n_{AA}(t, x) \tag{8}$$

$$n_{Aa}(t + 1, x + 1) = V_{Aa} W_{Aa} n_{Aa}(t, x) \tag{9}$$

$$n_{aa}(t + 1, x + 1) = V_{aa} W_{aa} n_{aa}(t, x). \tag{10}$$

Again, we focus on the proportion, p, of A alleles in the population, which is defined as

$$p(t, x) = \frac{n_{AA}(t, x) + \frac{1}{2} n_{Aa}(t, x)}{n(t, x)} \tag{11}$$

and at time $t + 1$ and patch $x + 1$ follows

$$p(t + 1, x + 1) = \frac{V_{AA} W_{AA} n_{AA}(t, x) + \frac{1}{2} V_{Aa} W_{Aa} n_{Aa}(t, x)}{V_{AA} W_{AA} n_{AA}(t, x) + V_{Aa} W_{Aa} n_{Aa}(t, x) + V_{aa} W_{aa} n_{aa}(t, x)}. \tag{12}$$

If we assume Hardy-Weinberg equilibrium and define $q = 1 - p$, then $n_{AA} = np^2, n_{Aa} = 2npq$, and $n_{aa} = nq^2$. Many well-known departures from this assumption are possible, with the magnitude of these departures influenced by the magnitude of differences among genotypes with respect to V and W. Nonetheless, these departures are, in practice, often small and have a tendency to vanish as time progresses. Assuming that departures from Hardy-Weinberg equilibrium are modest, we can approximate the dynamics of p recursively according to

$$p(t + 1, x + 1) \approx \frac{V_{AA} W_{AA} p^2(t, x) + V_{Aa} W_{Aa} p(t, x) q(t, x)}{\overline{VW}}, \tag{13}$$

where $\overline{VW} = V_{AA} W_{AA} p^2(t, x) + 2 V_{Aa} W_{Aa} p(t, x) q(t, x) + V_{aa} W_{aa} q^2(t, x)$ is the *mean spatiotemporal fitness* of the population. Again, this model is similar to the standard diploid model of selection (eqn. 2), except that fitness is now explicitly spatiotemporal. Also as before, in the event that there is no spatial fitness differential ($V_{AA} = V_{Aa} = V_{aa}$), the model reduces to the standard diploid model (eqn. 2).

This model has three equilibria. The equilibria at $p = 0$ and $p = 1$ are stable when $\frac{V_{aa}}{V_{Aa}} > \frac{W_{Aa}}{W_{aa}}$ and $\frac{V_{AA}}{V_{Aa}} > \frac{W_{Aa}}{W_{AA}}$, respectively. An equilibrium allowing for coexistence of the two alleles occurs at $p = \frac{V_{Aa} W_{Aa} - V_{aa} W_{aa}}{2 V_{Aa} W_{Aa} - V_{AA} W_{AA} - V_{aa} W_{aa}}$ and is stable when both $V_{Aa} W_{Aa} > V_{AA} W_{AA}$ and $V_{Aa} W_{Aa} > V_{aa} W_{aa}$. It is noteworthy that, unlike the haploid case, the vanguard can be populated by a stable mixture of A and a alleles but that this requires

some manner of trade-off between the life-history and dispersal traits of AA and aa genotypes that result in the Aa genotype having the highest spatiotemporal fitness.

Conditions under which spatial fitness matters

Now that we have established the notion of spatial fitness, we can explore a more general spatial arrangement to observe the conditions under which spatial fitness might change allele frequencies. Let us now imagine a two-patch system. Here, we have two patches ($x = 0$ and $x = 1$) occupied by n_0 and n_1 individuals, respectively. Otherwise, all other details are identical to the models developed above.

Haploid two-patch model

In this situation, following reproduction and dispersal, the number of individuals carrying each allele in patch 1 is:

$$n_{1,A}(t+1) = (1 - V_A)W_A n_{1,A}(t) + V_A W_A n_{0,A}(t) \tag{14}$$
$$n_{1,a}(t+1) = (1 - V_a)W_a n_{1,a}(t) + V_a W_a n_{0,a}(t). \tag{15}$$

Following identical steps to those above for the haploid model, dropping the notation but remembering all our n and p terms are variables in time, and recalling that $n_{0,A} = n_0 p_0$, $n_{0,a} = n_0(1 - p_0)$, and so on for patch 1. The proportion of A alleles in patch 1 can be expressed as:

$$p_1(t+1) = \frac{V_A W_A (n_0 p_0 - n_1 p_1) + W_A n_1 p_1}{V_A W_A (n_0 p_0 - n_1 p_1) + W_A n_1 p_1 + V_a W_a [n_0(1 - p_0) - n_1(1 - p_1)] + W_a n_1(1 - p_1)}. \tag{16}$$

It is clear from the numerator that the proportion of A alleles in the population depends upon standard fitness effects plus a spatiotemporal fitness term weighted by the allele-specific density gradient between the patches. To focus on the spatial fitness component, we now set $W_A = W_a$; $V_a = m$; and $V_A = m + s$ where $0 \le m \le 1$ is the spatial fitness (migration rate) of the a allele, and $0 < s \le 1 - m$ is the increase in spatial fitness associated with V_A. We then find that the proportion of allele A in patch 1 is:

$$p_1(t+1) = \frac{(m + s)(n_0 p_0 - n_1 p_1) + n_1 p_1}{(m + s)(n_0 p_0 - n_1 p_1) + n_1 p_1 + m[n_0(1 - p_0) - n_1(1 - p_1)] + n_1}. \tag{17}$$

The first term in the numerator represents the contribution to allele frequency from migration, and neatly separates the component due to standard gene flow, $m(n_0 p_0 - n_1 p_1)$, and spatial sorting, $s(n_0 p_0 - n_1 p_1)$. Spatial sorting plays no role only when $s(n_0 p_0 - n_1 p_1) = 0$, which occurs when $\frac{n_0}{n_1} = \frac{p_1}{p_0}$. Thus, unless the density gradient between our patches is perfectly balanced by a countergradient in allele frequecy, spatial sorting plays a role in changing the allele frequency. This is the same condition required for gene flow alone to contribute change to the allele frequency.

Diploid two-patch model

Similar reasoning can be applied to derive a map for the haploid two-patch model, focussing on spatial fitness. When we assume Hardy-Weinberg equilibrium within each patch following reproduction:

$$p_1(t+1) = \frac{n_1 p_1 [(1-V_{AA})W_{AA}p_1 + \frac{1}{2}(1-V_{Aa})W_{Aa}(1-p_1)] + n_0 p_0 [V_{AA}W_{AA}p_0 + \frac{1}{2}(V_{Aa}W_{Aa}(1-p_0))]}{n_1(1-\overline{V}_1)\overline{W}_1 + n_0 \overline{VW}_0},$$

(18)

where $\overline{V}_x, \overline{W}_x$ are mean spatial and temporal fitness in each patch.

We now need to make assumptions about dominance. In the simplest case, we assume complete dominance such that $V_{aa} = m$; $V_{Aa} = V_{AA} = m + s$. In this case (and again assuming no variation in fitness: $W_{AA} = W_{Aa} = W_{aa}$),

$$p_1(t+1) = \frac{n_1 p_1(1-(m+s))(p_1+1) + n_0 p_0(m+s)(p_0+1)}{2[n_1(1-\overline{V}_1) + n_0 \overline{V}_0]}.$$

(19)

Again, we attempt to separate effects due to migration and those due to spatial sorting. Here we face one condition for the numerator, and another condition in the denominator. In the numerator the effect of s is irrelevant when $n_1 p_1(p_1 + 1) = n_0 p_0(p_0 + 1)$; equivalently (when $n_1 > 0$), $\frac{n_0}{n_1} = \frac{p_1^2 + p_1}{p_0^2 + p_0}$. In the denominator, we find that $\frac{n_0}{n_1} = \frac{p_1^2 + 1 - p_1}{p_0^2 + 1 - p_0}$. Both these conditions must be met for spatial sorting to have no effect: $\frac{n_0}{n_1} = \frac{p_1^2 + p_1}{p_0^2 + p_0} = \frac{p_1^2 + 1 - p_1}{p_0^2 + 1 - p_0}$. And while these expressions are more complex than in the simple haploid case, they clearly represent a countergradient between density and allele frequencies.

Discussion

By focusing only on the evolutionary dynamics of an invasion front, the invasion front model developed here gives natural selection's shy younger sibling, spatial sorting, nowhere to hide (Shine et al. 2011). In the special context of the invasion front, spatial sorting is no longer a transient phenomenon – it is persistent through time – and this allows us to see that spatial sorting is akin to natural selection. Whereas natural selection operates to filter genotypes through time, spatial sorting operates to filter genotypes through space. Whereas natural selection filters genotypes on the basis of reproductive rate, spatial sorting filters genotypes on the basis of dispersal rate. As a result, we are justified in thinking of genotypes having both temporal and spatial aspects to their fitness.

In the special case of an invasion front, spatial sorting happens every generation and so exerts influence similar to that exerted by traditional natural selection. On invasion fronts, it is clear that natural selection and spatial sorting interact strongly; a situation most clearly seen in the stability criteria of the models here. The stability criteria in all cases express an imbalance between relative spatial and temporal aspects of fitness. If the relative spatial fitness of allele A is greater than the relative temporal fitness of allele a, then A will increase in frequency on the invasion front, even if this entails a reduction in traditional (temporal) fitness. Interestingly, our finding that A increases in frequency on an invasion front when $V_A W_A > V_a W_a$ is similar to the condition for invasion of the front recently derived by Deforet *et al.* (2017) using a reaction diffusion framework and competing clonal lines. Thus, it appears that spatial sorting is a directed evolutionary process that, in particular circumstances, can act to reduce the (traditional, temporal) fitness of a population. Under more complicated models, interactions between the temporal and spatial components of fitness become even richer, as revealed in a model of cane toad spread that pointed to temporal fitness giving rise to a boost in spatiotemporal fitness under certain circumstances (Perkins et al. 2013).

The invasion front is, however, a special case. We have simplified the spatial problem by making it unidirectional, like time. Because we focus on only the newly colonized patch, we only have one source of immigrants and have no need to account for emigration. While this approach does allow us to observe spatial sorting at play, it also makes our model quite specific. Our two-patch model removes the unidirectionality condition and revealed that spatial sorting will occur whenever there is a density gradient, unless that density gradient is anulled by a countergradient in allele frequency. That is, in almost any real (nonequilibrium)

setting. As with the invasion model, the two-patch model also shows that the role of spatial sorting is secondary to that of natural selection, operating only through an interaction with traditional (temporal) fitness.

A more general spatial setting would predict outcomes across all space, but this is a substantially more complex undertaking that must either be addressed numerically (e.g., Perkins et al. 2016) or analytically through the use of difficult mathematics (e.g., Benichou et al. 2012). It is our intention here to avoid such complexity, admittedly at some cost to generality. The intent of the formulation developed here is to go back to basics and integrate spatial sorting into basic evolutionary theory. The simple theoretical tools that we make use of are not new, of course. Population genetic models accounting for interactions between natural selection and gene flow go back decades (e.g., Nagylaki 1992) and have furthermore yielded great insight for applied topics, such as adaptation in response to a shifting environment (e.g., Case and Taper 2000). Perhaps by virtue of being a more recent conceptual development, the development of mathematical theory of spatial sorting has mostly bypassed simple models such as ours in favor of complex simulation models. Although belated, we hope that this paper will serve as a clear and incisive entry point to a literature that is already rich in biological detail and steeped in demonstration of applied value.

On the whole, our results speak to the striking similarity of action between natural selection and spatial sorting on an invasion front. The analogy we strike is very similar to that of Slatkin and Excoffier (2012) who show that serial founder events on an invasion front are the spatial analogue of drift. In both cases we see space playing the role of time in classic evolutionary models. This similarity hints that many theoretical results—built off the analogous standard models that underlie much of evolutionary theory—may also apply on invasion fronts subject to spatial sorting alone. That is, if there is no difference between alleles in aspects of temporal fitness (W terms are all equal), then the equations become identical to the standard haploid and diploid models, but they refer to spatial sorting rather than natural selection. Thus, we might expect that many theoretical results extending the basic model—exploring issues of dominance, frequency dependence, sexual conflict, drift vs selection, and so on—can be re-derived for spatial sorting. In particular, by combining Slatkin and Excoffier's (Slatkin and Excoffier 2012) approach to serial foundering with our notion of spatial sorting, it should now be possible to capture the drift/selection dynamics of dispersal-modifying alleles on invasion fronts. More generally, these simple models of selection and drift on invasion fronts offer a useful testbed for developing ideas about evolution in non-equilibrium populations before incorporating them into more powerful analytical frameworks (e.g., Benichou et al. 2012). At the same time, full elaboration of this body of theory will require bridging a subtle but potentially important technical distinction between our model (i.e., its pertinence to the most recently colonized location) and traveling wave models (i.e., their focus on the furthest location above which some threshold density is achieved). Even so, our simple formulation has the potential to provide a path towards fertile theoretical ground for years to come.

Acknowledgements

The manuscript was improved by thoughtful comments from Stephan Peischl, Susan Frances Bailey, and Frédéric Guillaume. BLP was supported by an Australian Research Council Future Fellowship.

References

Benichou, O., V. Calvez, N. Meunier, and R. Voituriez. 2012. Front acceleration by dynamic selection in Fisher population waves. Physical Review E 86:041908.

Bouin, E., and V. Calvez. 2014. Travelling waves for the cane toads equation with bounded traits. Nonlinearity 27:2233.

Bouin, E., V. Calvez, N. Meunier, S. Mirrahimi, B. Perthame, G. Raoul, and R. Voituriez. 2012. Invasion fronts with variable motility: Phenotype selection, spatial sorting and wave acceleration. Comptes Rendus

Mathematique 350:761–766.

Burton, O. J., J. M. J. Travis, and B. L. Phillips. 2010. Trade-offs and the evolution of life-histories during range expansion. Ecology Letters 13:1210–1220.

Case, T. J., and M. L. Taper. 2000. Interspecific competition, environmental gradients, gene flow, and the coevolution of species' borders. American Naturalist 155:583–605.

Crow, J. F., and M. Kimura. 1970. An introduction to population genetics theory. Burgess Publishing Company, Minneapolis, Minnesota.

Cwynar, L. C., and G. M. MacDonald. 1987. Geographical variation of lodgepole pine in relation to population history. The American Naturalist 129:463.

Deforet, M., C. Carmona-Fontaine, K. S. Korolev, and J. B. Xavier. 2017. A simple rule for the evolution of fast dispersal at the edge of expanding populations. arXiv:1711.07955 [q-bio].

Ditmarsch, D. van, K. E. Boyle, H. Sakhtah, J. E. Oyler, C. D. Nadell, É. Déziel, L. E. Dietrich, and J. B. Xavier. 2013. Convergent Evolution of Hyperswarming Leads to Impaired Biofilm Formation in Pathogenic Bacteria. Cell Reports 4:697–708.

Ellner, S. P., and S. J. Schreiber. 2012. Temporally variable dispersal and demography can accelerate the spread of invading species. Theoretical Population Biology 82:283–298.

Fisher, R. A. 1937. The wave advance of advantageous genes. Annals of Eugenics 7:355–369.

Fronhofer, E. A., and F. Altermatt. 2015. Eco-evolutionary feedbacks during experimental range expansions. Nature Communications 6:6844.

Haldane, J. B. S. 1924. A mathematical theory of natural and artificial selection—I. Transactions of the Cambridge Philosophical Society 23:19–41.

Hartl, D. L., A. G. Clark, and A. G. Clark. 1997. Principles of population genetics. Sinauer associates Sunderland.

Hughes, C. L., C. Dytham, and J. K. Hill. 2003. Evolutionary trade-offs between reproduction and dispersal in populations at expanding range boundaries. Proceedings of the Royal Society Biological Sciences Series B 270:S147–S150.

Lombaert, E., A. Estoup, B. Facon, B. Joubard, J. Grégoire, A. Jannin, A. Blin, and T. Guillemaud. 2014. Rapid increase in dispersal during range expansion in the invasive ladybird Harmonia axyridis. Journal of evolutionary biology 27:508–517.

Nagylaki, T. 1992. Introduction to Theoretical Population Genetics. Springer.

Ochocki, B. M., and T. E. X. Miller. 2017. Rapid evolution of dispersal ability makes biological invasions faster and more variable. Nature Communications 8:ncomms14315.

Olivieri, I., Y. Michalakis, and P.-H. Gouyon. 1995. Metapopulation Genetics and the Evolution of Dispersal. The American Naturalist 146:202–228.

Otto, S. P., and T. Day. 2007. A biologist's guide to mathematical modeling in ecology and evolution. Princeton University Press.

Perkins, S. E., C. Boettiger, and B. L. Phillips. 2016. After the games are over: Life-history trade-offs drive dispersal attenuation following range expansion. Ecology and Evolution 6:6425–6434.

Perkins, T. A., B. L. Phillips, M. L. Baskett, and A. Hastings. 2013. Evolution of dispersal and life-history interact to drive accelerating spread of an invasive species. Ecology Letters 16:1079–1087.

Phillips, B. L., G. P. Brown, and R. Shine. 2010. Evolutionarily accelerated invasions: The rate of dispersal evolves upwards during range advance of cane toads. Journal of Evolutionary Biology 23:2595–2601.

Phillips, B. L., G. P. Brown, J. M. J. Travis, and R. Shine. 2008. Reid's paradox revisited: The evolution of

dispersal in range-shifting populations. The American Naturalist 172:S34–S48.

Shine, R., G. P. Brown, and B. L. Phillips. 2011. An evolutionary process that assembles phenotypes through space rather than through time. Proceedings of the National Academy of Sciences 108:5708–5711.

Simmons, A. D., and C. D. Thomas. 2004. Changes in dispersal during species' range expansions. American Naturalist 164:378–395.

Slatkin, M., and L. Excoffier. 2012. Serial founder effects during range expansion: A spatial analog of genetic drift. Genetics 191:171–81.

Travis, J. M. J., and C. Dytham. 2002. Dispersal evolution during invasions. Evolutionary Ecology Research 4:1119–1129.

Van Valen, L. 1971. Group selection and the evolution of dispersal. Evolution 25:591–598.

Weiss-Lehman, C., R. A. Hufbauer, and B. A. Melbourne. 2017. Rapid trait evolution drives increased speed and variance in experimental range expansions. Nature Communications 8:ncomms14303.

Wright, S. 1931. Evolution in Mendelian populations. Genetics 16:97–159.