The impact of global selection on local adaptation and reproductive isolation

Gertjan Bisschop*[1], Derek Setter[1], Marina Rafajlović[2], Stuart J.E. Baird **[3], and Konrad Lohse **[1]

[1]Institute of Evolutionary Biology, University of Edinburgh, UK
[2]Department of Marine Sciences, Centre for Marine Evolutionary Biology, University of Gothenburg, Gothenburg, Sweden
[3]Institute of Vertebrate Biology, Academy of Sciences of the Czech Republic, Brno, Czech Republic

* correspondence to g.bisschop@sms.ed.ac.uk
** contributed equally

Abstract

Despite the homogenising effect of strong gene flow between two populations, adaptation under symmetric divergent selection pressures results in partial reproductive isolation: adaptive substitutions act as local barriers to gene flow, and if divergent selection continues unimpeded, result in complete reproductive isolation of the two populations, i.e. speciation. However, a key issue in framing the process of speciation as a tension between local adaptation and the

homogenising force of gene flow is that the mutation process is blind to changes in the environment and therefore tends to limit adaptation. Here we investigate how globally beneficial mutations (GBMs) affect divergent local adaptation and reproductive isolation. When phenotypic divergence is finite, we show that the presence of GBMs limits local adaptation, generating a persistent genetic load at the loci which contribute to the trait under divergent selection and reducing genome-wide divergence. Furthermore, we show that while GBMs cannot prohibit the process of continuous differentiation, they induce a substantial delay in the genome-wide shutdown of gene flow. In both cases, the strength of these effects depends crucially on two parameters: the rate of beneficial mutations relative to the rate of recombination events and the relative frequency of GBMs.

Introduction

Biologists tend to suffer from a cognitive bias (sensu Kahneman and Tversky (1972)) regarding adaptation: our education involves many examples of adaptation that astound us, and which can only be understood given an appreciation of deep time – where ephemeral humans have another strong cognitive bias. Given our inherent bias in favour of the power of adaptation, questions regarding the limits to adaptation must be addressed with every care designed into the scientific approach, if we wish objective answers. Felsenstein (1981) was among the first to appreciate this. He pointed out that by any measure there are many more niches than species and demonstrated that gene flow is a strong homogenising force that will tend to prevent populations adapting to different niches. As an illustration, each broad-leaved tree provides two equal-sized niches (patches) for feeding caterpillars: under-leaf and over-leaf. These niches likely differ in predator pressure. Why then does each species not bifurcate into an

under- and over-leaf phenotype? It seems they cannot, even though variation exists in nature for under- and over-leaf caterpillar lifestyles. Felsenstein's answer is that the divergent selection caused by differing predator pressure on the pair of patches is insufficiently strong to overcome the homogenising effects of gene flow between incipient patch-populations. His objective measure of the homogenising effects of gene flow became part of the canon of evolutionary biology and to some extent founded the study of limits to adaptation (Barton and Partridge, 2000; Lenormand, 2002).

Several verbal models and simulation studies have framed the process of speciation as a tension between local adaptation and the homogenising force of gene flow (Wu, 2001; Flaxman et al., 2013; Yeaman and Whitlock, 2011; Rafajlović et al., 2016). Flaxman et al. (2013) simulated divergent selection acting on many loci between two populations connected by gene flow and concluded that this model (intended as an extension of Felsenstein's) has an emergent property: populations can adapt to the opposing environmental stresses experienced in a pair of equal-sized patches even in the face of strong gene flow and, in the process, become (partly) reproductively isolated. Multilocus emergent properties had been described previously (e.g. Barton, 1983), and so, on superficial inspection, these results seem plausible. The model simulated by Flaxman et al. (2013) assumes that local adaptation is never-ending, i.e. each new mutation confers a fixed selective advantage locally. In contrast, other simulation studies of a pair of populations connected by gene flow have viewed local adaptation to a new set of fixed local optima (Yeaman and Whitlock, 2011; Rafajlović et al., 2016). However, while these population and quantitative genetic studies of divergent selection make different assumptions about the genetic basis of local adaptation (an eternal stream of local sweeps vs adaptation to a fixed set of local optima), they share an important feature: only locally beneficial mutations (LBMs) that

affect fitness as a function of the two patches are considered. This "adaptationist" simplification ignores a central tenet of Darwinian evolution: namely, that the mutational process is blind to changes in the environment (Barton and Partridge, 2000) and therefore will tend to limit adaptation.

Even if we ignore deleterious mutations, a substantial fraction of new beneficial mutations must be advantageous in many environmental contexts. Returning to our toy example of caterpillars in over and under-leaf patches, such globally beneficial mutations (GBMs) include all variants that increase caterpillar fitness in both leaf patches as well as any mutation that increases fitness at the adult stage. Even in the presence of a barrier to gene flow, such globally beneficial mutations will tend to selectively sweep across patches (Piálek and Barton, 1997). Previous analytic work has focused on the interaction between GBMs and LBMs in the context of adaptive introgression: how likely are GBMs to introgress from one species into another if they are linked to alleles with locally deleterious effects (Piálek and Barton, 1997; Hartfield and Otto, 2011; Uecker et al., 2015). Uecker et al. (2015) show that the probability of a single locally deleterious allele hitchhiking to fixation decays to zero over a distance $\frac{1}{2N_e r}$, where N_e is the effective population size and r is the scaled recombination rate. For example, given parameters for modern humans ($N_e = 10,000$ and $r = 10^{-8}$), the relevant distance is $\approx$ 5kb, shorter than the average gene. Given that many sexual species have higher scaled recombination rates, recombination should be frequent enough to prevent the majority of locally deleterious mutations from hitchhiking to fixation (Hartfield and Otto, 2011). While the sweep dynamics of individual variants that are linked to locally deleterious mutations have been characterized in some detail (Hartfield and Otto, 2011; Uecker et al., 2015), the flip-side of this interaction has received little attention: to what extent is local adaptation impeded when there is a constant supply of both GBMs

and LBMs arising in a genome? Moreover, the hitchhiking-fixation of locally deleterious mutations is not the main driving factor potentially slowing down adaptation: Sweep interference is, and occurs only while mutations are not fixed (Hill and Robertson, 1966; Barton, 1995).

Here we use simulations to investigate how the presence of GBMs affect the process of divergent local adaptation and reproductive isolation. We extend the existing simulation frameworks of (Yeaman and Whitlock, 2011; Rafajlović et al., 2016; Flaxman et al., 2013) in which divergence evolves under a constant high rate of mutational influx. In these models, local adaptation involves many loci and the dynamics resulting from the selective interference of locally and globally beneficial mutations cannot be captured by the analytic results that are available for the simpler case of a single introgressing locus.

Specifically, we i) ask to what extent adaptation to locally divergent trait optima is impeded by selective sweep interference from GBMs, ii) test how the effect of globally beneficial mutations depends on the assumptions about local adaptation (population genetics vs quantitative trait genetics models) and iii) consider how GBMs influence the evolution of reproductive isolation during the late stages of speciation.

Methods

General outline of the model of local adaptation

We study the impact of global selection by adding globally beneficial mutations (GBMs) to a multilocus model of divergent local adaptation similar to that studied by Yeaman and Whitlock (2011) and Rafajlović et al. (2016). Simulations were implemented in `SLiM3.3` (Messer et al., 2016). We consider two Wright-Fisher populations (discrete non-overlapping generations) of N_e diploid

individuals that exchange, on average, $M = 4N_e m$ migrants per generation and that experience soft selection. Recombination is modelled by assuming that cross-over events occur uniformly at random, i.e. we ignore gene conversion and physical constraints on double cross-over events. Both populations start out perfectly adapted to a shared local optimum set at 0 until the onset of divergent selection. An instantaneous change in the environment shifts their optima to θ_+, and $\theta_- = -\theta_+$, respectively. We thus assume that adaptation is *de novo* and that the optima are stationary. We further assume additive fitness effects, i.e. the phenotype of each individual is determined by the sum of effect sizes of all mutations affecting the trait under divergent selection. As in Yeaman and Whitlock (2011) and Rafajlović et al. (2016), the fitness of an individual in deme i with phenotype z is given by

$$W = e^{-\sigma(z-\theta_i)^2}.$$

New offspring is generated each generation until the preset fixed population size N_e is reached. Each individual's chance of being the parent to each of the new offspring is determined by their fitness W. σ affects the width of the fitness curve with smaller values representing weak selection (see 1).

Locally beneficial mutations (LBMs) arising at rate u_l per generation, per base pair, are assumed to be codominant with effect sizes drawn from a mirrored exponential distribution with mean μ_l.

Mutations can occur with equal probability at each position along a contiguous chromosome of map length l. We assume a large number of evenly spaced sites $(100,000)$, such that the probability of several mutations occurring at the same site becomes low (approaching the assumptions of an infinite sites model). If a site does get hit by a second mutation of the same type (LBM or GBM), this mutation will erase the previous one (house of cards model). This differs

from the model considered by Yeaman and Whitlock (2011) and Rafajlović et al. (2016) which assumes that mutations occur at a small number (between 50 and 2000) of evenly spaced loci with effects accumulating at loci that may be hit by several mutations (continuum-of-alleles model). By adjusting σ as well as μ_l (see table 1) we can adjust the number of mutations needed to fix in a population to reach the optimum, without changing their distribution of fitness effects (DFE), which is important for comparability once we add GBMs.

Adding globally beneficial mutations

We consider a scenario in which the two populations are also adapting to a shared moving optimum on a second orthogonal trait. However, we make two key modeling assumptions. The first assumption is that there is no pleiotropy, i.e. each beneficial mutation contributes either to local adaptation or global adaptation, but not both. This ensures that any detected interactions which limit local adaptation are only due to genetic hitchhiking (rather than pleiotropy). The second and more subtle assumption concerns the distribution of effect sizes for the GBMs. For a constantly moving optimum, the population will lag behind, and the adaptive walk results in a distribution dominated by intermediate-effect alleles (Kopp and Hermisson, 2009). To make GBMs comparable in effect size to LBMs, rather than simulating a moving optimum explicitly, we assume that the population trails the optimum by a substantial and fixed distanced. The fitness effect distribution of GBMs is modeled as an exponential distribution with mean s_g chosen to match the effect size distribution of LBMs at the onset of divergent selection, and GBMs arise at rate u_g. Combining the effects of LBMs and GBMs, an individual's relative fitness is defined as

$$W = e^{-\sigma(z-\theta_i)^2} \prod_{j}^{l} (1 + h_j s_j)$$

Parameter	Explanation	Value
N_e	Population size per patch	10000
M	Scaled migration rate	1
θ	Optimal phenotype for trait under divergent selection	$\theta_+ = 2$ $\theta_- = -2$
σ	Strength of selection	0.0025 (weak selection) 0.0125 (weak selection, polygenic) 0.025 (strong)
μ_l	Mean effect size mutations affecting local adaptation	0.05, 0.01 (polygenic)
l	Total number of sites	$1e5$
s_g	Mean exponential distribution selection coefficients GBMs	0.001
u/r	Ratio rate of beneficial mutations to recombination rate	0.001, 0.0035, 0.01, 0.1
T	Total duration of the run in $2N_e$ generations	5, 10 (polygenic)

Table 1: Model parameters and values used to simulate local adaptation in the presence of GBMs.

where $h = 0.5$ if an individual is heterozygous, and $h = 1$ if it is homozygous (i.e. we assume co-dominance).

To measure the impact of GBMs on trait divergence, we track the mean phenotype in both populations for the trait under local selection. Phenotypic divergence is measured as the difference Δz between population means, where 'perfect' differentiation corresponds to $\Delta z = 2\theta$. We also track the time and population of origin of LBMs and use the tree sequence recording in `SLiM3.3` to record and assess the distribution of cross-population coalescence times at the end of each run using `tskit` (Kelleher et al., 2018).

To facilitate comparisons between scenarios with and without GBMs, the rate u_l of LBMs is kept constant throughout. Since the relative ratio of GBMs to LBMs is unknown in nature we explore varying u_g/u_l, but ensure that the total rate of beneficial mutations ($u = u_g + u_l$) is biologically plausible given empirical estimates of *de novo* mutation rates and the distribution of fitness effects (see Choice of parameter space). For each parameter combination we ran 200 replicate simulations.

Globally beneficial mutations in the BU2S model

We have re-implemented the model studied by Flaxman et al. (2013) in *SLiM3* (Haller and Messer, 2019; Messer et al., 2016) by making minor changes to the model described above. For ease of comparison, we have kept all parameters (u_l, u_g, M, T) the same as in the phenotype-based model described above (Table 1). And as before, we assume that new mutations are either LBMs or GBMs. LBMs no longer affect fitness by means of an explicit phenotype. Instead they confer a homozygous fitness $1 + s_i$ in population 1 but $1/(1 + s_i)$ in population 2. As before co-dominance is assumed. Selection coefficients are drawn from an exponential distribution with mean $s_l = 0.001 = s_g$.

Choice of parameter space

A key quantity is the ratio μ/r which determines the build up of the mutational load relative to the speed at which recombination can decouple GBMs introduced by migration from their locally deleterious background.

We have sought to cover a biologically plausible range of parameter space including humans and *Drosophila*. For example, assuming that a fraction $p_a = 0.01$ of mutations is selectively advantageous (Schneider et al., 2011) and given estimates of the per base rates of recombination and *de novo* mutation in *Drosophila melanogaster* of $r = 10^{-8}$ (Chan et al., 2012) and $\mu = 3.5 \cdot 10^{-9}$; Keightley et al. (2009)), we expect $u/r = 0.0035$ for beneficial mutations. The analogous ratio in humans is 0.01 (assuming $\mu/r \approx 1$). Given that we ignore recombination heterogeneity along the chromosome and because regions with decreased recombination rate (e.g. telomeres) may experience increased interference, we extended our parameter space to higher and lower u/r. Simulations were run for two different lengths of sequence $lr = 0.5$ Morgan and 0.05 Morgan to test this ratio argument.

Results and Discussion

The results section is structured as follows: we first describe the effect of GBMs on adaptation to a fixed set of local trait optima. We then ask how GBMs affect genome-wide divergence measured in terms of the distribution of between population coalescence times. Finally, we investigate GBMs in the context of the more extreme model of continued divergence considered by Flaxman et al. (2013) and discuss their effect on the evolution of reproductive isolation in the late stages of speciation.

GBMs reduce local adaptation

Comparing the mean divergence along the local trait axis between simulations with and without GBMs shows clearly that GBMs reduce local adaptation (figures 1, S1). Unsurprisingly, the effect of GBMs on local adaptation depends on both the relative strength of selection on LBMs and GBMs (s_g/s_l) and their relative mutation rates (u_g/u_l). In particular, GBMs have a negligible effect in the strong selection regime, i.e. when selection on LBMs is much stronger (10x) than on GBMs at the onset of local adaptations (figure S3).

Rather than slowing down local adaptation, GBMs halt the process at a certain point, thus decreasing the (maximum) possible trait divergence between populations. In other words, GBMs induce a constant genetic load (Lenormand, 2002) through a reduction in the proportion of LBMs segregating in the population which promote divergence and an increase in the proportion of those which serve to diminish it (figures 3 and S5). This causes the trait divergence to assymptote to an equilibrium value below the population optima, and this equilibrium is dynamic, i.e. subject to high variance both within and among individual runs (figures S1, 2).

We have assumed the effect of GBMs depends on the ratio of beneficial

mutations to recombination u/r. We tested this by scaling up the mutation and recombination rate 10 fold (by simulating a 10 times greater length of sequence) both for the polygenic and the non-polygenic case. Unsurprisingly, scaling up only holds in the polygenic case (figures S2 bottom row). In the non-polygenic case increasing the mutational supply results in the fast build up of local adaptation limited to a small fraction of the genome and is little affected by sweeping GBMs which act uniformly across the entire sequence. The more loci required for perfect adaptation, the larger the chance that GBMs arise in close proximity. In what follows, we therefore focus our attention on the weak-selection, polygenic case.

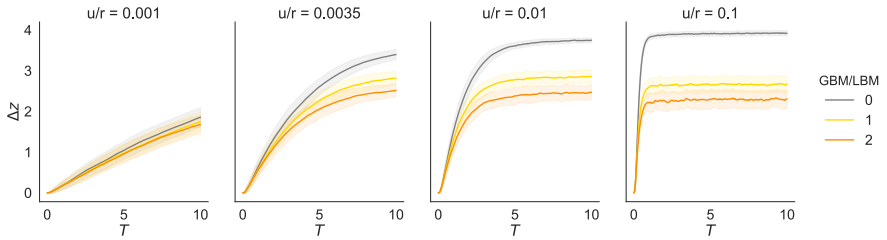

Figure 1: The time it takes for two populations to diverge in average phenotype (Δz) depends on u/r. GBMs reduce phenotypic divergence in the weak selection regime ($s_g = s_l$) with $\mu_l = 0.01$ (polygenic scenario). Gray line represents the scenario without GBMs, coloured lines with GBMs ($u_g/u_l \in [1, 2]$) Since the map length $R = 5cM$ is fixed, the mutational input and the number of GBMs sweeping at any given time increases from left to right. This explains why adaptation is faster and the impact of GBMs larger. The width of the shaded area is 2 standard deviations.

To understand why the observed dynamics differ from the impact of single introgression events of beneficial mutations (here GBMs) associated with a deleterious load (here LBMs), it is useful to relate the continuous process of sweeping GBMs to the adaptive trajectory of each population: at the onset of divergent selection, each population is displaced far from the new local optimum. As a consequence, under Fisher's geometric model, the initial local adaptation

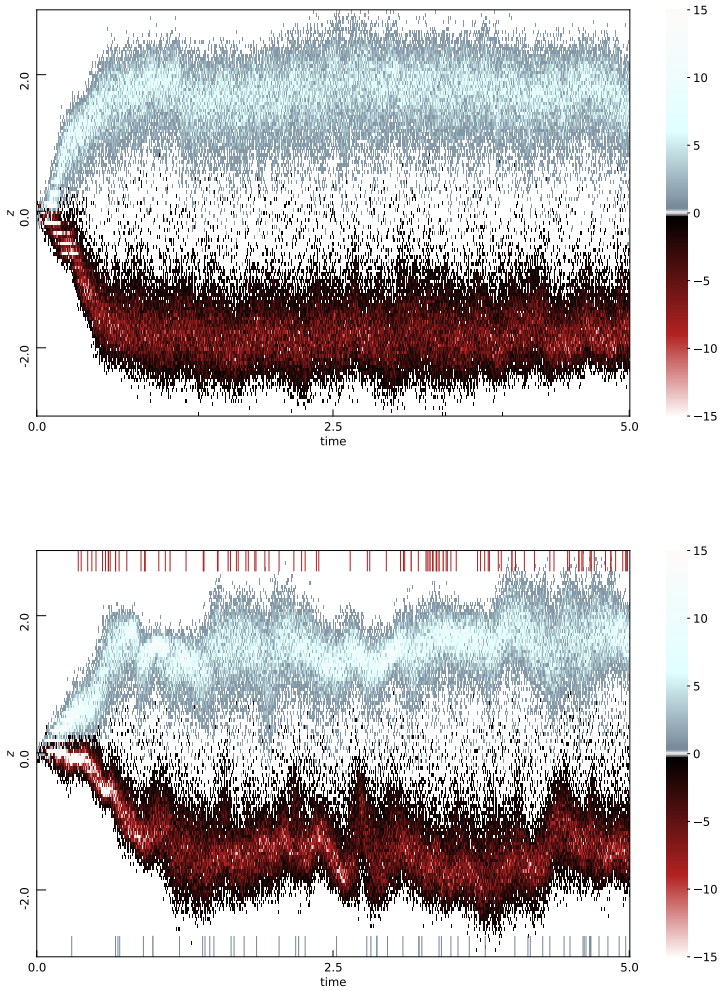

Figure 2: The evolution of individual trait values z for a single run without (above) and with (below) GBMs (for the weak selection, non-polygenic case). $u/r = 0.01$. Individuals (sample of 100 per population) are binned by phenotype, the shade of (red or blue depending on the population) each dot representing the number of individuals per bin. Lighter shades indicate a larger number of individuals. The times at which individual GBMs fix globally and their population of origin are demarcated along the top and bottom. Only GBMs with $s >= 0.0035$ are shown. Data corresponds to a single iteration from figure 1 panel 3.

is due to a few mutations of large effect, and during this rapid phase, GBMs are not likely to occur in close proximity. The subsequent approach to the new optimum proceeds through mutations of progressively smaller effect sizes (more densely distributed). Thus the probability that GBMs arise in close proximity is increased as is the chance of deleterious LBMs to rise to appreciable frequencies before they are dissociated from sweeping GBM by recombination (figures 2, S4). At the same time, due to the fact that the fitness landscape becomes less steep upon approaching the optimum, the ratio s_l/s_g decreases. This increases the predicted frequency at which a locally deleterious LBM will initially segregate upon recombining away from the sweeping GBM (Uecker et al., 2015). Once a LBM is at appreciable frequency, GBMs arising in the same population affect its trajectory (figure S4).

To summarise, the effects of sweeps on the adaptive trajectories of populations overlap for many generations. Analytical results for sweeps occurring in isolation from each other (single introgression events) will not capture the potential sweep interference (between LBMs and GBMs) arising. For example, even though there is only a very limited fraction of the genome around an established LBM along which GBMs might affect the frequency trajectory of that LBM in the 'wrong' population, these introgression events happen over and over again, with each migrant that makes it across. Even if the deleterious LBM recombines away and is purged, its sojourn time in the 'wrong' deme will be increased relative to a scenario without GBMs. Sweep interference can occur throughout the time overlap of these polymorphic sojourns, and so its potential to limit adaptation is increased with GBMs. Additionally, unlikely events of a GBM and a deleterious LBM fixing together will accumulate (figure 3 right bottom pane) which are compensated by new LBMs. This turn-over is visible in the age distribution of LBMs which in the presence of GBMs shows a skew

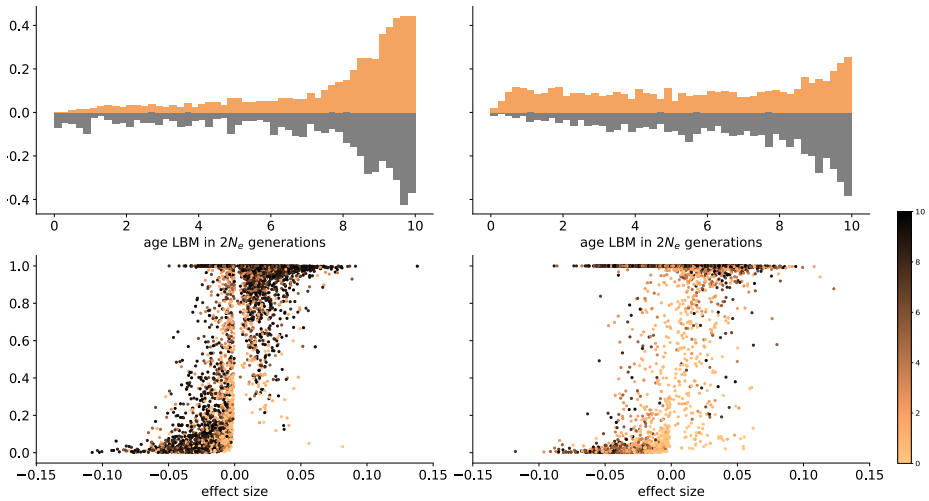

Figure 3: Top: the age distribution of LBMs (across the genome and across 50 replicate runs, weighted by their frequency and effect size) for the case of weak selection and polygenic adaptation. Age is rescaled to $2N_e$ units with 0 representing the end of the simulation run. Top histogram represents LBMs with a positive effect size, bottom histogram LBMs with negative effect size. Bottom row: LBM effect size plotted against the allele frequency (in the population adapting to θ_+) in the dynamic equilibrium phase ($10 \times 2N_e$ generations). Young and old LBMs are shown in orange and black respectively (see color bar). Results without GBMs (left) and with GBMs (right).

towards the recent past (figure 3 top) while in the absence of GBMs, LBMs contributing to local adaptation tend to be old, i.e. have fixed during the initial phase of local adaptation (figure 3 left).

GBMs reduce genome wide divergence

The evolution of reproductive isolation can be characterized as a gradual reduction in effective migration rate (m_e) (Barton and Bengtsson, 1986). A measurable (in practice) metric for m_e is the genome-wide distribution of pairwise coalescence times between species, $f(T_2)$. As long as gene flow is ongoing, re-

cent coalescence is likely. However, once speciation is complete, $f(T_2)$ is fixed and can only shift pastwards. Inspection of $f(T_2)$ genome-wide (averaged across simulation replicates) shows that in the absence of GBMs, selection against maladapted migrants reduces m_e (Aeschbacher et al., 2017) and hence the chance of recent coalescence compared to the neutral expectation (figure 4). In contrast, the continued global sweeps induced by introgressing GBMs have the opposite effect of increasing m_e and the fraction of genome with recent ancestry.

Booker et al. (2019) have recently shown that GBMs also present a challenge for attempts to infer the targets of divergent selection in genome scans and may be impossible to distinguish from LBMs using F_{st} as a diagnostic summary statistic (Charlesworth, 1998; Cruickshank and Hahn, 2014). However, given the different effects LBMs and GBMs have on the distribution of between population coalescence times, it should be possible to distinguish between them using richer summaries of sequence variation.

GBMs delay the build up of strong reproductive isolation

We have so far only considered the effect of GBMs on adaptation to a fixed set of local optima. The scenario of recurrent sweeps on LBMs considered by Flaxman et al. (2013) is both simpler and more extreme (and arguably less realistic). While Flaxman et al. (2013) do not consider phenotypes explicitly, their model assumes that local selection is never-ending which is equivalent to assuming that populations adapt to eternally diverging local optima. In this case, strong reproductive isolation ("genome congealing") is inevitable, however GBMs still affect the speed at which such reproductive isolation evolves. Although GBMs only result in a modest increase in m_e when measured over a single generation as the relative average fitness of migrants (figure S7), their long-term cumulative effect is to delay the completion of strong reproductive isolation. The delay in

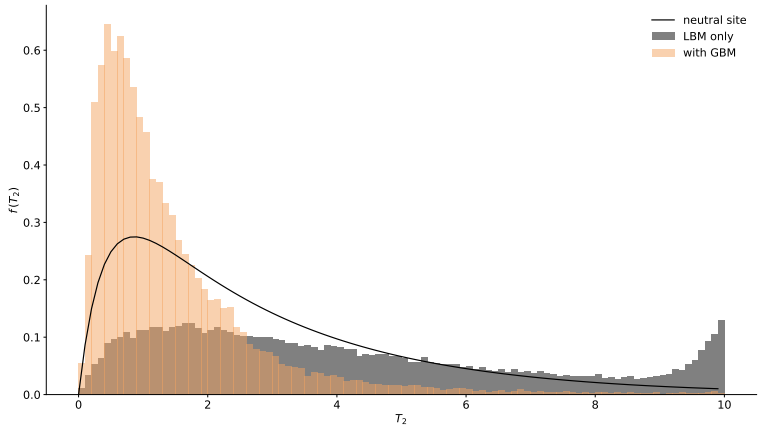

Figure 4: The genome-wide distribution of pairwise coalescence times $f(T_2)$ for the case of weak selection and polygenic adaptation. The neutral expectation (assuming $M = 1$) (Lohse et al., 2016, eq. 10) is shown as a gray solid line, distributions (across 50 simulation replicates) in the presence and absence of GBMs are shown in gray and red respectively. 69% (without GBMs) and 99% (with GBMs) of pairwise coalescence times are smaller then $10T$. Data corresponds to figure 1 panel 3.

the genome wide shutdown of gene flow results in a much wider distribution of between population coalescence times (figure S8).

Conclusions and future directions

We set out to address an 'adaptationist' imbalance in existing models of evolution under divergent selection: where evolution's stream of mutations, some of which give rise to better organisms, is redirected on the advent of a pair of patches, such that better organism mutations (GBMs) stop and only better LBMs continue. These imbalanced models have a very specific mutational process: general conclusions drawn are therefore extrapolation. We asked, for similar models, what happens when GBMs don't stop.

We have shown that GBMs can limit local adaptation in the face of gene flow and may impede the evolution of reproductive isolation regardless of whether we consider local adaptation to a fixed set of optima or as a runaway process. The effect of GBMs depends crucially on two parameters: the ratio of beneficial mutation to recombination events u/r (which determines the interference between selection on GBMs and LBMs) and the relative frequency of GBMs (u_g/u_l). We have only considered u_g/u_l of up to 2. While the ratio of GBMs to LBMs is an unknown quantity in nature which depends critically on how far real populations lag behind global optima, it seems plausible that there are many more mutations that affect global fitness than local adaptation suggesting results should be developed for a range of u_g/u_l ratios.

Our intent was to explore a more general mutational model for adaptation under divergent selection. However, for the sake of comparability, we have followed many of the same simplifying assumptions as existing work, i.e. complete symmetry both in terms of patch-specific demographic parameters (N_e, M) and selection (an instantaneous switch to symmetric trait optima). We have also

introduced our own symmetric simplification: the same DFE for GBMs and LBMs. Given that such many-fold symmetry is biologically unrealistic, an important task for future work is to test the robustness of these results to violations of symmetry. If our results are specific to symmetries, then by our own argument, any general conclusions drawn are extrapolation. While we can venture to make educated guesses about relaxing symmetry assumptions for individual parameters (e.g. GBMs should reduce local adaptation in both populations even when only one is given a new optimum), the effect of GBMs under a completely general model is difficult to intuit, especially if one were to allow for hard as well as soft selection. It appears general models, where unintended extrapolation is no longer a risk, are urgently needed as many areas of endeavour could be affected. For example: How do GBMs affect selection on recombination modifiers? Would we expect stronger selection on inversions that link LBMs (Navarro and Barton, 2003)?

Acknowledgements

We thank Nick Barton, Jarrod Hadfield for discussions, Tom Booker for constructive criticism of the manuscript. This work was supported by an ERC starting grant (ModelGenomLand). KL was supported by a NERC fellowship (NE/L011522/1). SJEB was supported by GACR grant 17-25320S. MR was supported by the Hasselblad Grant for Female Scientists. This work has made use of the resources provided by the Edinburgh Compute and Data Facility (ECDF) (http://www.ecdf.ed.ac.uk/).

References

Aeschbacher, S., Selby, J. P., Willis, J. H., and Coop, G. (2017). Population-genomic inference of the strength and timing of selection against gene flow. *Proceedings of the National Academy of Sciences*, 114(27):7061–7066.

Barton, N. H. (1983). Multilocus Clines. *Evolution*, 37(3):454.

Barton, N. H. (1995). Linkage and the limits to natural selection. *Genetics*, 140(2):821–841.

Barton, N. H. and Bengtsson, B. O. (1986). The barrier to genetic exchange between hybridising populations. *Heredity*, 57(3):357–376.

Barton, N. H. and Partridge, L. (2000). Limits to natural selection. *BioEssays*, 22(12):1075–1084.

Booker, T. R., Yeaman, S., and Whitlock, M. C. (2019). Global adaptation confounds the search for local adaptation.

Chan, A. H., Jenkins, P. A., and Song, Y. S. (2012). Genome-Wide Fine-Scale Recombination Rate Variation in Drosophila melanogaster. *PLoS Genetics*, 8(12).

Charlesworth, B. (1998). Measures of divergence between populations and the effect of forces that reduce variability. *Molecular Biology and Evolution*, 15(5):538–543.

Cruickshank, T. E. and Hahn, M. W. (2014). Reanalysis suggests that genomic islands of speciation are due to reduced diversity, not reduced gene flow. *Molecular Ecology*, 23(13):3133–3157.

Felsenstein, J. (1981). Skepticism Towards Santa Rosalia, or Why are There so Few Kinds of Animals? *Evolution*, 35(1):124–138.

Flaxman, S. M., Feder, J. L., and Nosil, P. (2013). Genetic hitchhiking and the dynamic buildup of genomic divergence during speciation with gene flow. *Evolution*, 67(9):2577–2591.

Haller, B. C. and Messer, P. W. (2019). SLiM 3: Forward Genetic Simulations Beyond the Wright-Fisher Model. *Molecular Biology and Evolution*, 36(3):632–637.

Hartfield, M. and Otto, S. P. (2011). Recombination and hitchhiking of deleterious alleles. *Evolution*, 65(9):2421–2434.

Hill, W. G. and Robertson, A. (1966). The effects of linkage and the limits to artificial selection. *Genetical Research*, 8(1966):269–294.

Kahneman, D. and Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454.

Keightley, P. D., Trivedi, U., Thomson, M., Oliver, F., Kumar, S., and Blaxter, M. L. (2009). Analysis of the genome sequences of three Drosophila melanogaster spontaneous mutation accumulation lines. *Genome Research*, 19(7):1195–1201.

Kelleher, J., Thornton, K. R., Ashander, J., and Ralph, P. L. (2018). Efficient pedigree recording for fast population genetics simulation. *PLoS Computational Biology*, 14(11):1–21.

Kopp, M. and Hermisson, J. (2009). The genetic basis of phenotypic adaptation II: The distribution of adaptive substitutions in the moving optimum model. *Genetics*, 183(4):1453–1476.

Lenormand, T. (2002). Gene flow and the limits to natural selection. *Trends in Ecology and Evolution*, 17(4):183–189.

Lohse, K., Chmelik, M., Martin, S. H., and Barton, N. H. (2016). Efficient strategies for calculating blockwise likelihoods under the coalescent. *Genetics*, 202(2):775–786.

Messer, P. W., Ellner, S. P., and Hairston, N. G. (2016). Can Population Genetics Adapt to Rapid Evolution?

Navarro, A. and Barton, N. H. (2003). Chromosomal Speciation and Molecular Divergence–Accelerated Evolution in Rearranged Chromosomes. *Science*, 300(5617):321–324.

Piálek, J. and Barton, N. H. (1997). The spread of an advantageous allele across a barrier: the effects of random drift and selection against heterozygotes. *Genetics*, 145(2):493–504.

Rafajlović, M., Emanuelsson, A., Johannesson, K., Butlin, R. K., and Mehlig, B. (2016). A universal mechanism generating clusters of differentiated loci during divergence-with-migration. *Evolution*, 70(7):1609–1621.

Schneider, A., Charlesworth, B., Eyre-Walker, A., and Keightley, P. D. (2011). A method for inferring the rate of occurrence and fitness effects of advantageous mutations. *Genetics*, 189(4):1427–1437.

Uecker, H., Setter, D., and Hermisson, J. (2015). Adaptive gene introgression after secondary contact. *Journal of Mathematical Biology*, 70(7):1523–1580.

Wu, C. I. (2001). The genic view of the process of speciation. *Journal of Evolutionary Biology*, 14(6):851–865.

Yeaman, S. and Whitlock, M. C. (2011). The genetic architecture of adaptation under migration-selection balance. *Evolution*, 65(7):1897–1911.

Supplementary information

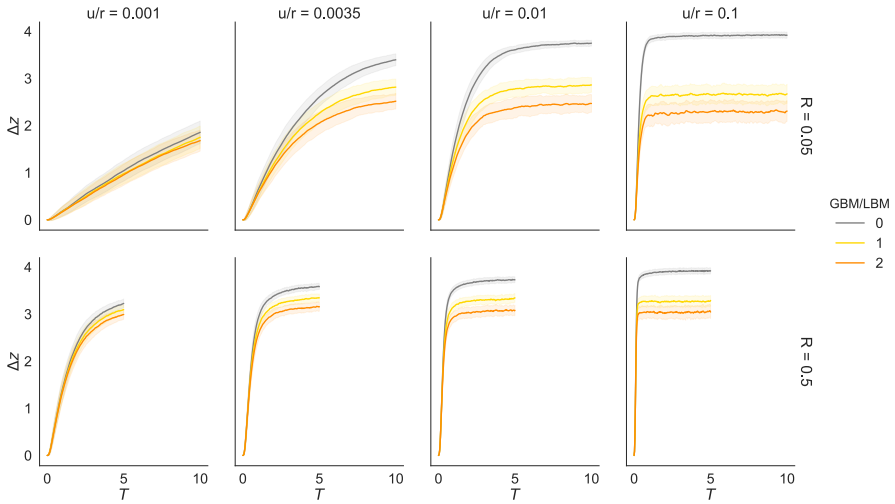

[h]

Figure S1: The time it takes two populations to diverge in average phenotype (Δz) for the polygenic weak selection case. Top row $R = 5cM$, bottom row $R = 50cM$.

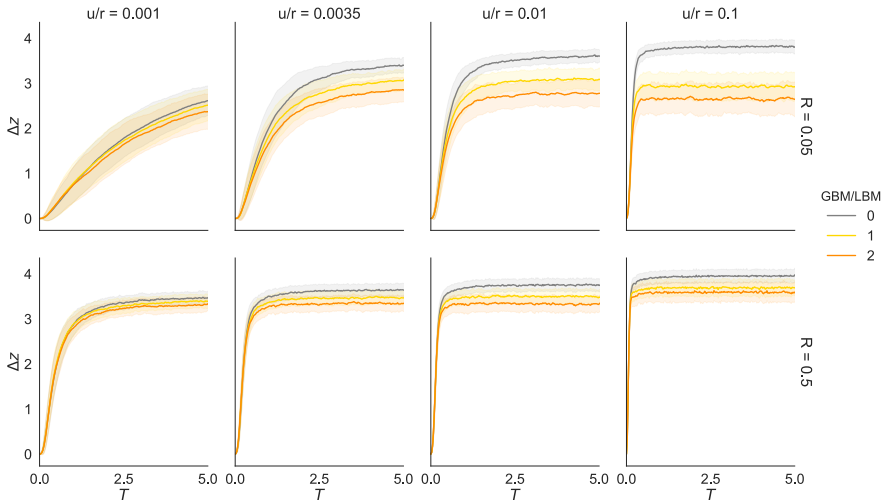

Figure S2: Impact of scaling: The time it takes two populations to diverge in average phenotype (Δz). GBMs reduce phenotypic divergence in the non-polygenic weak selection regime ($s_g = s_l$). Top row $R = 5cM$, bottom row $R = 50cM$.

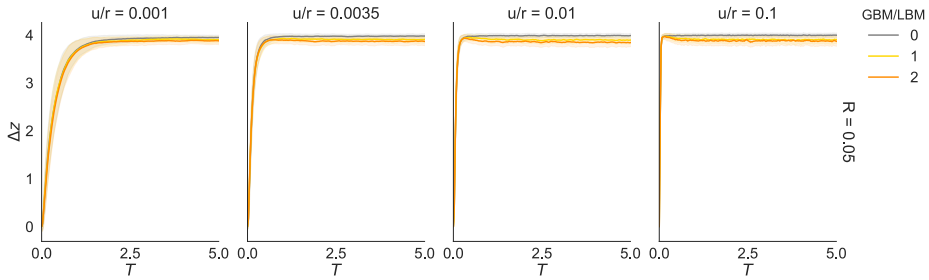

Figure S3: The time it takes for two populations to diverge in average phenotype (Δz) depends on u/r. GBMs don't affect phenotypic divergence in the non-polygenic strong selection regime ($s_g = 10 \times s_l$ at $t = 0$) with $\mu_l = 0.05$. Maplength $R = 5cM$.

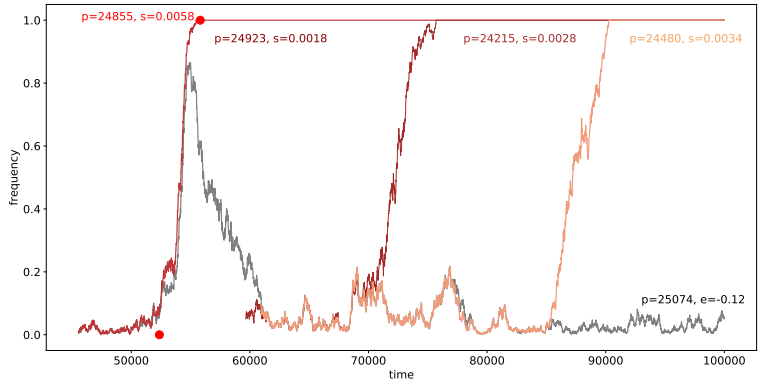

Figure S4: Frequency trajectories of three GBMs in one population (dark red, red and orange) for the weak selection scenario with $u/r = 0.01, \mu_l = 0.05$ (polygenic scenario). A single maladaptive LBM is tracked as well. At first the dark red GBM manages to recombine away from the LBM, but a new GBM (red dot) arises on that 'original' background. The deleterious LBM hitchhikes to high frequency. After recombining two more new GBMs affect is trajectory. The LBM does not fix however. This data set corresponds to panel (1,3) of figure S1

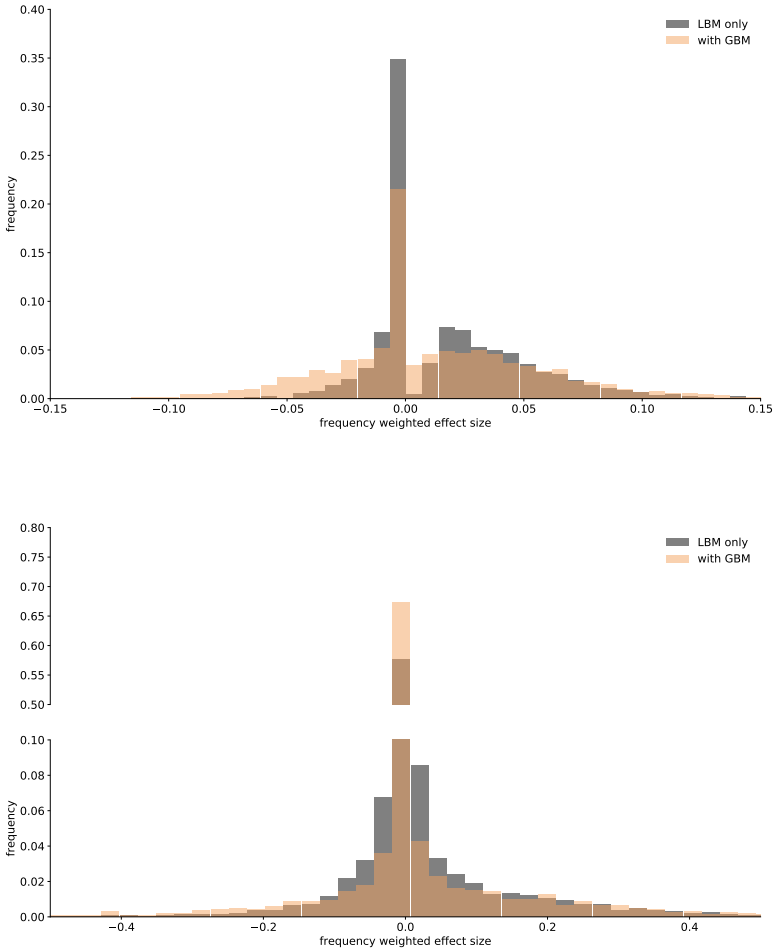

Figure S5: The distribution of frequency-weighted effects of LBMs in the dynamic equilibrium phase (at $10 \times 2N_e$ generations) for the polygenic case (top) and for the non-polygenic case (at $5 \times 2N_e$ generations, bottom), with (orange) and without (gray) GBMs. The top plot corresponds to figure S1 panel (1,3); the right plot, S2 panel (1,3)

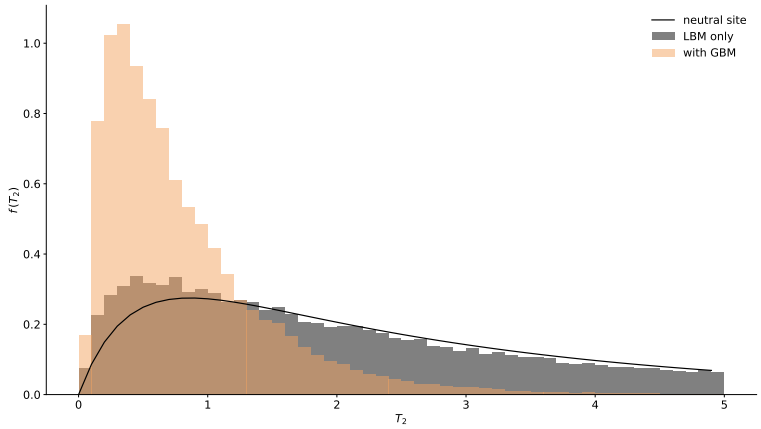

Figure S6: The genome-wide distribution of between population coalescence times $f(T_2)$. The neutral expectation (assuming $M = 1$) (Lohse et al., 2016, eq. 10) is shown as a gray solid line, distributions (across 200 simulation replicates for the weak selection non-polygenic case, $u/r = 0.01, \mu_l = 0.05$) in the absence and presence of GBMs are shown in gray and orange respectively. 86% and 99% of coalescence times respectively are smaller then $5T$. This data corresponds to figure S2 panel (1,3)

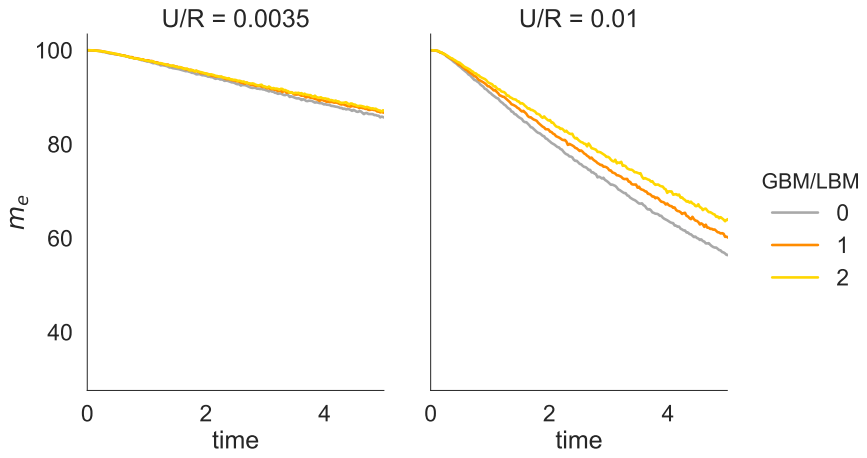

Figure S7: The relative fitness of migrants for the weak selection scenario $s_l = s_g = 0.001$ in the BU2S model (Flaxman et al., 2013). Time measured in $2N_e$ generations, the ratio of GBMs/LBMs goes from 0 (gray) to 2 (orange).

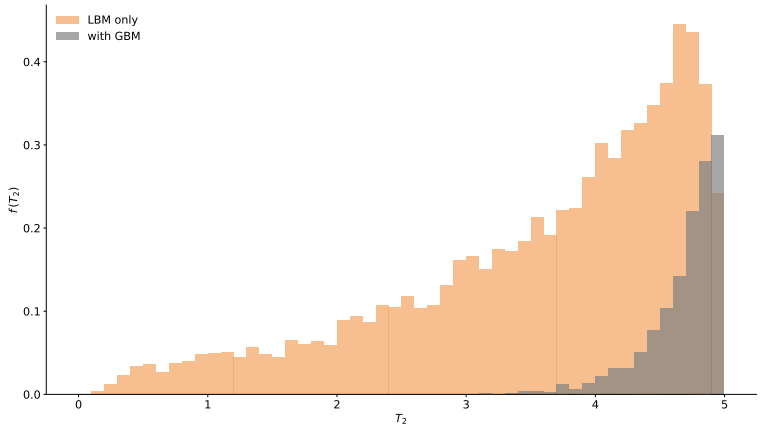

Figure S8: The genome-wide distribution of pairwise coalescence times $f(T_2)$ across 50 replicates for the BU2S model ($u/r = 0.01, s_g = s_l = 0.001, M = 1$). Without GBMs the genome wide migration shutdown occurs around 4 $\times 2N_e$ generations in the past. In the presence of GBMs the two populations share ancestry more recently at a substantial fraction of genome. 13% (LBMs only) and 73% (with GBMs) of coalescence times are smaller then $5T$