Dear Recommender,

We thank you and the reviewers for considering our manuscript for Peer Community in Evolutionary Biology. We addressed the comments of the reviewers in this revised version, which we believe to be improved. In particular, we aimed at improving two major points:

1) we clarified our methodology to better reflect that there are two modeling steps: (1) a new sequentially Markov coalescent model, which is fitted to the full genome; and (2) a causal model of the population genetic mechanisms shaping genome-wide diversity (represented as a linear model), which is fitted to window-averaged quantities. Importantly, these window-based quantities are not free parameters (as one reviewer pointed out, this would be a methodological step-back), but posterior averages resulting from an empirical Bayesian approach using the model fitted during step 1.
2) we enhanced the description of how our study compares with previous works, both in terms of methods and results. We note that no previous approach aimed at jointly inferring the landscapes of recombination, mutation and genealogies. The most advanced approaches incorporate mutation rate variation using a proxy, and report the total variance in diversity explained after removing the effect of mutation rate variation. How much of the total diversity is explained by the different causal factors has not been explored previously, and it is a novel result we bring in this study.

Below we list all points raised by the reviewers. Our detailed answers are typeset in italics.

Reviewer 1:

R1.1 The authors present an inference model that expands on their previously published iSMC model, that infers how nucleotide diversity in the genome is jointly influenced by genealogical history, recombination, and large-scale regional spatial variation in mutation rates. Applying this model to Drosophila, they argue that variation in mutation rates is the primary driver of nucleotide diversity at varying spatial resolutions throughout the genome. This has important implications for various inference tasks in population genomics. Overall, I found the approach innovative and interesting and so the work presents a valuable contribution to the field. At the same time, I remain skeptical about the biological conclusions about Drosophila, in light of the stark difference from the long history of work arguing for the key importance of linked selection.

A1.1 We thank the reviewer for his positive assessment of our approach. We would like to stress that our approach does not deny the importance of linked selection. Instead, we argue that some of the most prominent publications in the field measure the importance of linked selection after discarding the effect of mutation rate variation. In the revised version of our manuscript, we provide more detailed comparison with previous studies and emphasize this point (see second paragraph of the Discussion).

R1.2 While the importance of mutation rate variation on different scales seems intuitive to me, I remain unsure whether the variance decomposition framework provides adequate support for the authors' claims.

A1.2 We answer this concern in two ways, which we detail below. Briefly, we first note that our modeling is done at the genome scale, and not in windows, a point that we better state in the revised version of the manuscript (see A1.5). Second, we added an alternative assessment of the relative importance of each factor in our linear models (see Reviewer 2's comment R2.1 and our answer A2.1), which gives additional support to our conclusions.

Major comments

R1.3 The authors model pi (pairwise nucleotide diversity at SNPs) as a response variable, linearly related to their estimates of several parameters: population scaled recombination rate, mutation rate and coalescent time and the interaction of coalescent time with mutation rate (a term for which the justification is not entirely clear).

*A1.3 We added an explanation for the inclusion of an interaction term (**lines 221-223**): "Because of the interplay between genealogical and mutational variance, we tested the improvement that including an interaction term between θ and τ brought to the fit of the linear models."*
*We also added a justification for the linear model of pi itself (**lines 214-217**): "The justification for a linear model of π is that for sufficiently small genome-wide average diversity θ (a requirement which is met in Drosophila, as E[θ] ~1e-2) the per-site heterozigosity π=1-exp(-θ*τ) can be well approximated by θ*τ, the first term in the Taylor series expansion of 1-exp(-θ*τ)."*

R1.4 The parameters themselves are estimated in an iSMC framework on a small part of the genome. From the population genetic perspective, this seems like a step back compared to work in the last decade and a half to explain variation in nucleotide diversity along the genome. Of course, if prediction is king, then any approach that does the job well is merited. However: (a) It is hard to interpret results, even qualitative ones. For example, my sense is that the transition probabilities of 1 for recombination rate and mutation rate (lines 170-185)--- estimates that are at the limit of the possible range for the estimand—suggest that the analysis window size needs to be explored in more depth, and the model is ill-specified for 10kb windows.

A1.5 We estimate coalescent model parameters exclusively using the whole genome, so that all windows (which are only defined after fitting iSMC) share the same set of parameters, independently of their size. These parameters correspond to the demographic model (i.e. ancestral population sizes), parameters of the prior distributions for the recombination and mutation rates (as such, these are hyper-parameters), and transition probabilities between rate classes. The window-specific quantities are obtained by averaging in various window sizes the posterior decoding of the HMM. The empirical Bayesian posterior decoding methodology allows us to get site-specific estimates of quantities related to the hidden states while being economical in parameters. We also note that no transition probability is estimated to either 0 or 1. delta_rho ~ 0.99999 reflects that the model detects a change in recombination rate on average every 10 kb and delta_theta ~ 0.999999 reflects that the model detects a change in mutation rate on average every 100 kb, both of which seem to be reasonable estimates (more is discussed in the main text).
As the initial parameter estimation is highly time-consuming, we fit our model to a subset of the genome, the entire chromosome 2L arm. We then use the inferred demography and prior rate distributions when decoding all chromosome arms. We have no a priori reason to believe that the distribution of recombination rates, mutation rates, and TMRCA differs among autosomes. Still, in any case, the posterior decoding is robust to the inferred HMM parameters (in this case, using the distributions inferred from 2L to reconstruct dedicated landscapes for 2R, 3L and 3R). Supporting these claims, the chromosome-specific linear models yield consistent results (Supplemental Tables S6, S7 and S8).
We have rewritten the Methods session to make the presentation of the methodology more clear.

R1.6 I also do not see how this provides support for the claims on the relative importance of smaller-scale factors (e.g. nucleotide context) that are not modelled at all, to the best of my understanding, versus larger-scale factors (e.g. replication timing).

A1.7: Our approach only allows us to infer mutation rate along the genome, without making any assumptions regarding the underlying mechanisms that cause such variation. We discuss the

relative importance of mutation rate variation on genetic diversity, but our model does not enable us to disentangle distinct sources for this variation. To avoid any ambiguity, we rephrased ***line 185****, in the result section to "This suggests that our model mostly captures large-scale rather than fine-scale variation in the mutation rate."*
*In the discussion, we evoke the possible molecular mechanisms of the variation we observe, without arguing about their relative contribution (****lines 566-576****):*
"It will be interesting to test whether mutation events follow similar patterns, now that the impact of various sequence motifs on local μ is being more thoroughly investigated (DeWitt et al., 2021; Kim et al., 2021; Oman et al., 2022). Unraveling the factors that shape the mutation landscape at different genomic scales will likely provide important insight. For example, can the large-scale variation in mutation rates that we found in D. melanogaster be partially explained by aggregation of short (differentially mutable) sequence motifs, or is it driven by independent genomic features? As the molecular underpinnings of adaptive mutation landscapes become elucidated (e.g., what kind of proteins, epigenetic markers and sequence motifs are involved in increasing replication fidelity in functionally constrained regions and eventually decreasing it where polymorphism would tend to be beneficial) we will gain a better understanding of how flexible such phenotype is and how prevalent it is expected to be in different phylogenetic groups."

R1.7 (b) I think the bar has to be set high here, given the careful modeling done by Sattath et al. PLoS Genetics 2011 to consider effects of sweeps and demography while controlling for mutation rate variation (albeit imperfectly, by "normalizing" by divergence) and then Elyashiv et al., who extended the results to include (again, coalescent-based) effects of background selection and incorporated information about functional importance of individual sites and genomic regions. Here, in a variance decomposition framework, I can't help but wonder whether the reason the mutation rate factor explains the most variation is just that it is the most free to vary across adjecent windows, such that variance unaccounted for by the model can be easily absorbed in this term. (Here, simulations with little spatial mutation rate variation, and with or without local (e.g. nucleotide-specific, or just smaller than the analysis window) would be useful.)

A1.7 As pointed out by the reviewer, the aforementioned work used an external source of information to control for mutation rates. More precisely, the authors fix the local relative mutation rate in order to estimate the relative importance of background selection and selective sweeps. As such, they aim at explaining genetic diversity once the effect of mutation rate variation has been removed (Figure 2 in Elyashiv et al shows such "scaled" diversity. We further checked with one of the authors of this work that our interpretation of their result is indeed correct - Guy Sella, personal communication).
In contrast, our approach aims at estimating the "raw" genetic diversity by jointly accounting for mutation rate variations and TMRCA variation, which are previously jointly inferred by iSMC. We would further like to note that in our modeling framework, mutation rate variation is no more or less free to vary than the recombination rate, and not particularly more free to change than the TMRCA. As we pointed out, the posterior estimation of rates and TMRCA is independent of the chosen window size, as the model is fitted on the complete chromosome. In fact, the inferred parameters show that the recombination rate varies more often than the mutation rate (about 10x, see ***lines 183-185****). In our simulation study, we studied different scales of variation for the mutation rate (see Figure 5). We emphasize this point throughout the revised Methods section.*

R1.8 (c) Some comparisons with previous results of the landscape of mutation, recombination and selection in Drosophila is needed. It's one thing for the model to do well on simulations where the generative process is picked by the researchers; it is another to show that the variation plausibly reflects real biology. e.g. Assaf et al. 2017, recombination maps, and estimates of the intensity of

selection with Elyashiv et al. 2016.

*A1.8 The ability of iSMC to infer recombination rate was extensively benchmarked in a previous study (Barroso et al., 2019). In the revised version of this manuscript, we added a comparison with Comeron's empirical recombination map. We show that the correlations are 0.594 at the 50 kb scale, 0.693 at the 200 kb scale and 0.865 at the 1 Mb scale (**lines 193-195**). Assessing the accuracy of the mutation rate inference is more challenging, as a mutation map is not available for Drosophila. Using the divergence with D. yakuba as a rough proxy for the mutation rate (since the mutation map evolves over time), we report a correlation of 0.197 (**line 517**). Furthermore, we now directly compare our estimates with those of Elyashiv et al. (**lines 503-523**):*
"Besides Comeron, Elyashiv et al. (2016) also used patterns of nucleotide variation to fit models of linked selection along the fruit fly genome. Using substitution rates at synonymous sites as a proxy for local mutation rates, they employed their selection estimates to predict genome-wide diversity in windows from 1 kb to 1 Mb. Their models predict 44% (100 kb) and 76% (1 Mb) of the distribution of scaled nucleotide diversity in D. melanogaster. However, owing to the scaling that removes the effect of mutation rate variation, the percentages in the Elyashiv study represent the part of variance explained by linked selection once the effect of mutation rate variation has been discarded (see also Murphy et al., (2022) for a similar model). As such, the R^2 they report quantify the goodness-of-fit of different models of selection (e.g., background selection alone vs background selection + selective sweeps) instead of the importance of linked selection to π, and are, therefore, not directly comparable with our estimates. We note, however, that the remaining variance in their models may be due to mutation rate variation not grasped by synonymous divergence – an imperfect proxy for μ, either because of selection on codon usage or because the mutation landscape has evolved since the divergence of the two species. (Along the same vein, the correlation between our 50 kb mutation maps and genome-wide divergence between D. melanogaster and D. yakuba is only moderate, Spearman's rho = 0.197, p-value = 3e-09.) The differences between our approaches to capture linked selection are also worth discussing. While Elyashiv et al. (2016) relied on elaborate models of selection that embody strong assumptions, we leaned on the spatial distribution of τ, similarly to Palamara et al. (2018). As such, our approach has less parameters (11 compared to 36), rendering it less susceptible to mis-specifications of the selection model, which could be commonplace (e.g., the presence of epistasis and/or fluctuating fitness effects over time)."

Other comments

R1.9 Does increasing this granularity improve model fit? (Conversely, does using a model with fewer parameters improve performance in smaller window sizes?) See for example Michaelson et al., 2012,) using an HMM to model effects of regional mutation rate on de-novo mutation counts in human trios,

A1.9 As stated in A1.4, the iSMC model is fitted to the whole genome with few parameters. Compared to a PSMC model, iSMC adds only two parameters for the recombination rate variation and two parameters for the mutation rate (and reduces the number of demographic parameters by using a splines construction). The size of windows used in the linear modeling is independent of the number of parameters initially fitted by iSMC.

R1.10 The authors mention in the discussion (line 523-525) that it would be interesting to assess the extent to which large-scale variation in mutation rates is explained by the composition of sequence motifs (with different mutation rates) in each genomic window. This is more or less already a consensus in the human mutation rate literature (most recently, Fang et al. showed that motif-specific mutation rate estimates explain a significant portion of large-scale variation in mutation rates in humans). It would be straightforward and worth including some simple analyses to demonstrate the relevance of motif-specific mutation rates in the regional mutation trends. One idea

might be to identify particular 3-mers that are overrepresented in the windows with the highest and lowest genetic diversity. (the Drosophila mutation spectrum characterized by Assaf et al., 2017 would be useful here).

A1.10 We thank the reviewer for their suggestion. We edited the corresponding part in the discussion to better reflect the state of knowledge. Assessing whether our estimates of mutation rates correlate with particular sequence motifs in D. melanogaster *is an exciting perspective of our work, which we now mention* ***lines 570-572****.*

R1.11 Relatedly, I'm curious how the authors expect this model might generalize to mammalian species whose genomes are significantly influenced by the extreme hypermutability of methylated CpG sites. Since Drosophila is not affected by mCpG hypermutability, I'd expect the overall mutation landscape to look more "flat" than in mammals, so if applied to a mammalian species, is it possible that this model would simply recapitulate regional variation in mCpG content? Stratifying some of the analyses by mutation (sub)type might yield some interesting insights into how theta, rho, & tau are interacting to shape genome diversity.

A1.11 We thank the reviewer for raising this point. We find it difficult to make any prediction at this stage. In Drosophila, our model captured large scale variations of the mutation rate, and we do not know what is the finest variation scale that we can possibly capture, notably because it depends on the recombination rate landscape. Correlating the local mutation rate with sequence features is an interesting avenue for future research, and mammalian species a relevant model for doing so (besides their much larger genome size).

R1.12 Line 172-174: Does fitting on a single chromosome arm sufficiently model genome-wide mutation rate heterogeneity and/or accommodate large genomic regions that may harbor unusually high or low mutation rates in drosophila? An interesting counterexample might be chr8p in humans, which is known to have a much higher mutation rate than elsewhere in the genome.

*A1.12 We agree with the reviewer that this is a point of concern. In this study, we restricted our analyses to the four chromosome arms 2L, 2R and 3L, 3R, as they show comparable diversity patterns. In other datasets where this would not be the case, improved resolution may be obtained by estimating chromosome-specific parameters (while assuming a common demography). We further note that the posterior decoding is robust to the inferred distribution of parameters. Increasing the number of rate categories when discretizing the distributions should also enable a finer characterisation of local rates. However, we have not thoroughly tested this possibility in this study because of the computational resources demanded to run such models. We now also mention this computational bottleneck in the main text (****lines 161-170****).*

Minor comments

R1.13 Line 50-51: "new stab" is a bit vague—what exactly are these recent papers attempting to clarify/augment RE: Lewontin's Paradox?

A1.13 This sentence has been rephrased as: "The observation that π does not scale linearly with population size across species (Lewontin, 1974) was termed "Lewontin's Paradox", and recent work has taken a new stab at this old problem by modeling the effect of natural selection (Buffalo, 2021; Galtier and Rousselle, 2020)."

R1.14 Line 490: "according to their position in the genome" is a bit too reductive here, as it implies that mutation patterns in functional loci are purely a function of the surrounding genomic features—mechanisms like transcription-associated mutagenesis and transcription-coupled repair uniquely impact the mutation landscape of transcribed loci, independently of those loci's spatial positioning

in the genome.

A1.14 This is now rephrased as: "Nevertheless, the mutation landscape should play a role in the dynamics of natural selection by modulating the rate at which variation is input into genes (and other functionally important elements) depending on their position in the genome."

R1.15 Line 540-545: Venkat et al. is another good ref here, as they provide a compelling example of how a more detailed mutation model (in this case, one that accounts for multi-nucleotide mutations) can confound selection scans.

A1.15 We thank the reviewer for pointing out this reference, which we have included in the revised version.

R1.16 Fig 2. Please add labels on the left for the 50kb/200kb/1Mb scales
Fig 5. Legend: typo in sentence "in A and B, results are displayed according [to] simulated parameters"
Fig 6. Please make font sizes of panel labels larger
Fig 7. As in Fig 2, please add labels above each column of panels for the scale
Line 227: typo, simulatios -> simulations

A1.17 We thank the reviewer for pointing out these errors, which we have corrected as suggested.

Reviewer 2:

In this manuscript, Barroso & Dutheil propose an extension of a statistical method that jointly infers the genomic landscape of genealogies (or "local Ne"), recombination rates and mutation rates. They benchmark the method with simulations and apply it to Drosophila melanogaster from Zambia. They find that, at the genomic window lengths that they analyze (50Kb, 200Kb and 1Mb), the mutation landscape seems to be the most important determinant of the levels of genetic diversity along the Drosophila genome. This conclusion is somehow contradicting Comeron 2014, where he concluded that the genetic diversity landscape is mostly affected by linked selection (or tau, or TMRCA, using Barroso & Dutheil terminology). However, the authors do a good job of reasoning why both studies seem to reach contradictory conclusions. This manuscript is relevant to the population genomics community because it makes available a powerful tool and it brings back to the spot light the mutation landscape. I agree that the mutation landscape is often an overlooked ingredient to explain the genetic diversity landscape within a genome.
I've divided this revision into 4 sections.

1. Is the science sound, with a logical narrative and well-supported results and conclusions?
The manuscript follows a logical narrative and the methods are sound. The simulations and benchmarking are convincing. The literature context provided in the introduction and discussion is very helpful (below I suggest a couple of more papers to back up some ideas tho). However, some aspects require further clarification.

R2.1 I do not think that the partial R^2 is a good way to assess the relative importance of each variable (mutation rate, recombination rate and TMRCA) on the levels of genetic diversity. The standardized regression coefficients, which are the regression coefficients obtained from estimating a model on the standardized variables (mean = 0, standard deviation = 1), are better suited for this job IMO. I would also suggest reporting the variance inflation factors in Table 1.

A2.1 We agree with the reviewer that having standardized explanatory variables is a very useful alternative. We reran our linear models using this transformation and have updated the manuscript

tables to include the new coefficients). In all cases, the linear coefficient attributed to theta is higher than the one attributed to tau (by a factor of >3). We also included the variance inflation factors in Table 1, Supplemental Tables S2, S9, noting that in the main text, we focus on the GLS coefficients because this model corrects for heteroscedasticity and auto-correlation of the residuals. However, the linear coefficients may be affected by model violations (normality, homoscedasticity and independence of error terms), while R^2 is a measure of goodness of fit, which does not rely on such assumptions. As R^2 is interpretable in terms of "proportion of variance explained by the model/variable", we find it easier to compare with results of other studies (which have also discussed their results in these terms). Therefore, although we now report both linear coefficients from standardized variables and R^2 in the main text, we discuss our results mainly in terms of the latter.

R2.2 I wonder why the authors do not try to validate their mutation rate and recombination rate estimates using empirical measures. I understand that perhaps empirical measures of the mutation rate (using mutation accumulation inbreed lines?) might be hard to find, but the empirical recombination landscape from Comeron is publicly available.

A2.2: see response A1.8 to reviewer 1's comment.

R2.3 Moreover, how the global or genome-wide TMRCA inferred with the new method compares to previous demographic estimates in this population? If they are not similar then what can be the cause?

A2.3 Demographic inference was not a central point in our study, so we fitted a spline model to capture its effect with a minimum number of parameters. The inferred variation of ancestral population sizes, however, is consistent with previous reports, with evidence for a recent bottleneck. In ***lines 186-190****, we write:*
"Our inferred genome-wide average (0.036) is in line with previous estimates (Chan et al., 2012), and the coalescence rates (which, for the purpose of this paper, comprise a collection of nuisance parameters used to refine our estimates of τ, ρ and θ along the genome) suggest a relatively recent ~4-fold bottleneck followed by a partial recovery (Supplemental Figure S1)."

R2.4 Related to the previous point. For future implementations maybe in the Discussion. Could it be possible to plug-in empirical mutation and recombination landscapes to infer the local TMRCA (or "local Ne")?

A2.4 We thank the reviewer for this interesting suggestion. We note, however, that fixing the recombination and mutation landscapes would be an extension of the classical PSMC/MSMC. The iSMC framework, which models these landscapes as unknown latent variables, is not required for this purpose. Fixing the landscape would not add any parameter to the classical PSMC, as the emission and transition probabilities would be fully specified. However, while no additional parameter would be needed, the heterogeneity of the probabilities would impede most of the computational heuristics used to increase the speed of the likelihood calculation. To our knowledge, no software has yet implemented this possibility.
*Another possibility to infer "local Ne" is to extend the iSMC framework further to modulate Ne along the genome. This is, in principle, possible since Ne affects both the transition and emission probabilities, independently of the recombination and mutation rates. We added a sentence about the possibility of modulating the TMRCA by a landscape of Ne (****lines 523-524****). While we believe this to be an exciting perspective, heuristics will be needed to make this approach computationally feasible.*

R2.5 Regarding the DFE of Drosophila. The authors simulate a shape = 1 (row 654), this is

equivalent to an exponential distribution and will produce way more weakly deleterious mutations than the ones expected when the shape = 0.3-0.4 (which is the value more commonly estimated in the literature for this species, see https://academic.oup.com/mbe/ and others). This excess of weakly deleterious mutations could explain, I believe, the results explained from rows 406 to 415.

*A2.5: We thank the reviewer for raising this point. The shape of 1.0 was selected following the example from the SLIM manual (section 17.4 Detecting the "dip in diversity"). Intending to be conservative, we chose such parameters to maximize the effect of background selection and, therefore, the extent of model violation imposed on iSMC. We improved the discussion of these results in the revised manuscript (**lines 424-427**).*

R2.6 In here (https://www.ncbi.nlm.nih.gov/) Spencer et al. 2006 use "wavelet techniques to identify correlations acting at different scales" which in essence is very similar to Barroso & Dutheil work.

A2.6 We thank the reviewer for pointing out this work to us. The Wavelet methodology would indeed be relevant here, offering a better (although more complex) approach to the three selected scales of analysis. We will certainly consider it in future studies.

R2.7 here (https://onlinelibrary.wiley.com/doi/10.1002/bies.201200150) Martincorena and Luscombe 2012 review "the main forces driving the evolution of local mutation rates and identify the main limiting factors" which might be relevant for sentences in rows 495-500.

A2.7 We agree and have included the reference.

R2.8 Bear in mind that Castellano et al. 2018a was already published in GBE "long" ago. The cited preprinted version might be outdated.

A2.8 We updated the reference.

R2.9 Could the authors provide some intuition about the "five mutation rate classes, five recombination rate classes and 30 coalescent time intervals" used? Why this setting and not another one? How relevant is this choice in downstream analyses?

*A2.9: In our implementation, the number of classes is only a matter of discretisation. (In the original PSMC, the number of TMRCA classes is also linked to the number of epochs modeled in the demographic scenario, which is not the case in iSMC, which uses splines instead.) The number of parameters used is, therefore, independent of the choice of discretisation. The number of classes, however, determines the number of hidden states in the HMM (we have n = 5*5*30 = 750 states) as well as the memory used (scaling in O(n)) and computational time (scaling in O(n^2)). A higher number of classes generally leads to a more precise inference but uses more resources. We find that three classes already give good results, but since we specifically aimed at studying the recombination and mutation landscapes, we considered five categories in this work. We did not have the computational resources to increase further any of these numbers without reducing another and settled for this compromise. We clarified these aspects in the revised version of the manuscript, **lines 161-170**:*
"We note that the number of classes and time intervals do not affect the number of estimated parameters, as our implementation uses splines in place of a skyline model. The finer the discretization, the more precise the inference. However, the memory usage and the likelihood computation time scale linearly and quadratically with the total number of hidden states, respectively. Importantly, in the case of mutation and recombination rate distributions, the number of classes will directly impact the minimum and maximum values that the posterior estimates can

take. We selected five classes to obtain a reasonably good resolution. The total computation time for fitting the model to chromosome 2L was about 1 month on a high-performance cluster."

R2.10. Are there obscure passages that you (or a potential reader) can't go through?
3.1 I do not understand the first sentence of the paragraph starting at row 154.
3.2 In figures 4 and 5 it might be helpful to scale the y-axes in log units?
3.3 Typo at row 245 "the contribution of contribution"?

A2.10: We streamlined the sentence mentioned in 3.1 and fixed the typo, but we still prefer to show the y-axis of figures 4 and 5 in a linear scale.

R2.11. Potential extra analysis only if interesting enough to the recommender and/or author: Relevant to rows 346-355. I am just curious to know how much tau (or TMRCA) varies along the genome in the absence of selection compared to the presence of selection. Could the authors show some density plots in both scenarios? I think it is often assumed that in the absence of selection genetic diversity is entirely explained by the mutation landscape. Still, it seems that the TMRCA (and genetic diversity) can vary along the genome stochastically in the presence of recombination. This is an interesting finding that should be further highlighted.

A2.11 We performed an additional set of simulations identical to our set of simulations with selection, but setting the rate of deleterious mutations to 0. We then compute the average TMRCA in windows, as in the manuscript. By comparing the two sets, we observe that the mean TMRCA is reduced in the scenario with selection, but that the variance is indeed increased (note also that the variance in TMRCA decreases with window size even under neutrality, as we emphasize when discussing our results in the main text):

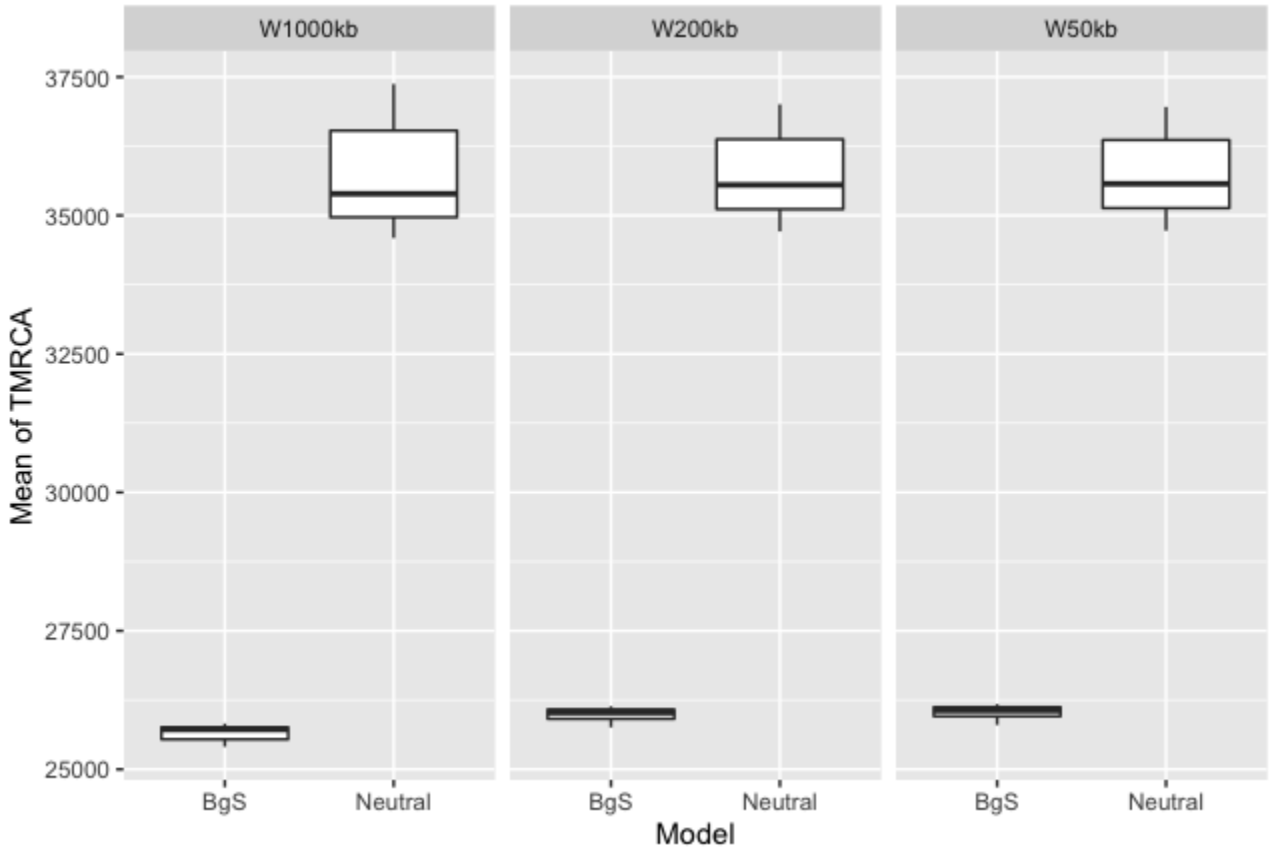

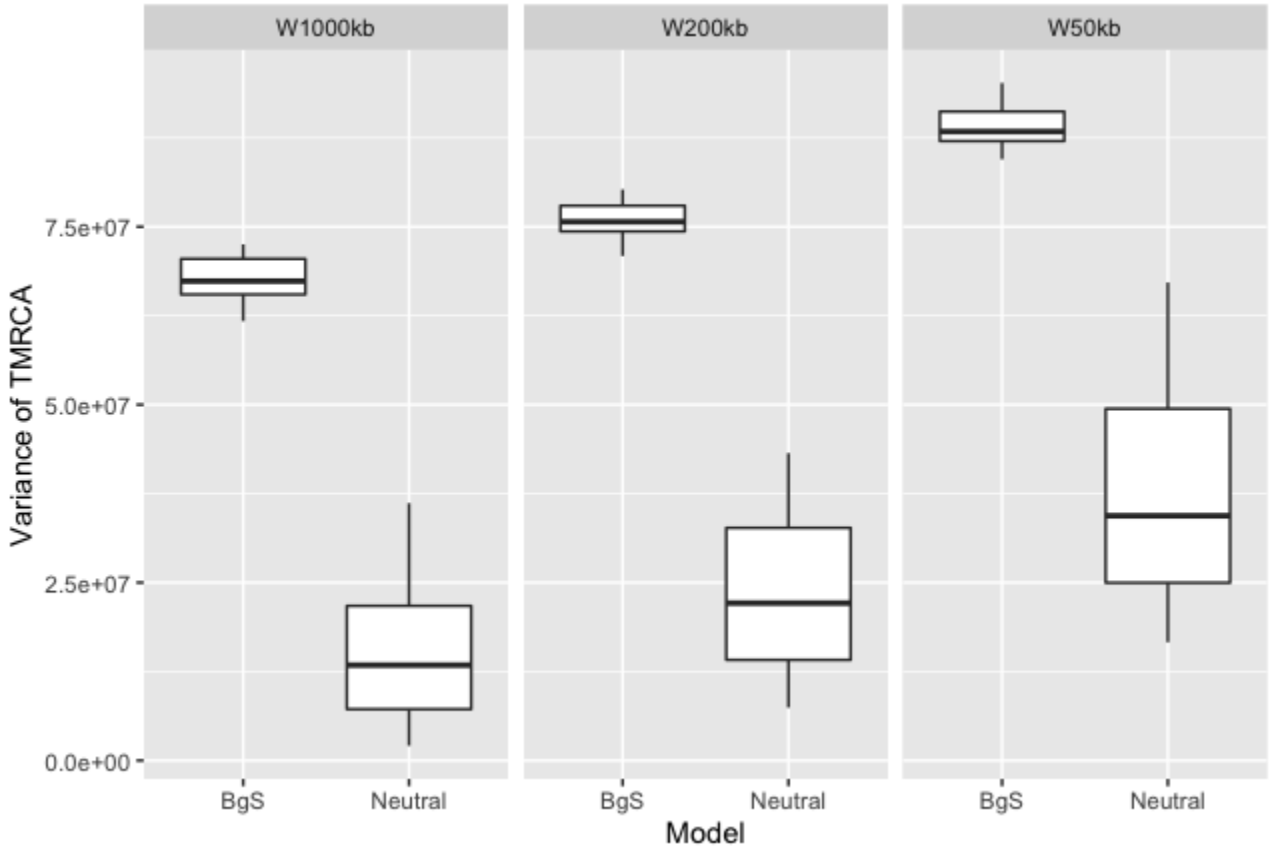

*We added these simulations and corresponding scripts to the supplementary data repository but would prefer not to add these figures to the manuscript, which already contains a broad set of simulation results. Regarding the stochastic variation in the TMRCA, we added a reference to a recent MBE paper by Schlichta et al. (2022) showing the effect of range expansions (demography) and recombination rates on throughs of diversity along the genome (**lines 373-376**). These are caused by variations in the TMRCA landscape, meaning that they are actually quite similar in spirit to the phenomena we report in our neutral simulation study (Figure 5).*