

Shifting Balance on a Static Mutation–Selection Landscape: A Novel Scenario of Positive Selection

Christopher T. Jones,^{*,1} Noor Youssef,² Edward Susko,^{1,3} and Joseph P. Bielawski^{1,2,3}

¹Department of Mathematics and Statistics, Dalhousie University, Halifax, Nova Scotia

²Department of Biology, Dalhousie University, Halifax, Nova Scotia

³Center for Comparative Genomics and Evolutionary Bioinformatics, Dalhousie University, Halifax, Nova Scotia

*Corresponding author: E-mail: cjones2@dal.ca

Associate editor: Claus Wilke

Abstract

A version of the mechanistic mutation–selection (MutSel) model that accounts for temporal dynamics at a site is presented. This is used to show that the rate ratio dN/dS at a site can be transiently >1 even when fitness coefficients are fixed or the fitness landscape is static. This occurs whenever a site drifts away from its fitness peak and is then forced back by selection, a process reminiscent of shifting balance. Shifting balance is strongest when the substitution process is not dominated by selection or drift, but admits interplay between the two. Under this condition, site-specific changes in dN/dS were inferred in 78–100% of trials, and positive selection (i.e., $dN/dS > 1$) in 10–40% of trials, when sequence alignments generated under MutSel were fitted to two popular phenomenological branch-site models. These results demonstrate that positive selection can occur without a change in fitness regime, and that this is detectable by branch-site models. In addition, MutSel is used to show that a site can be occupied by a sub-optimal amino acid for long periods on a fixed landscape when selection is stringent. This has implications for the interpretation of constant-but-different site patterns typically attributed to changes in fitness. Furthermore, a version of MutSel with episodic changes in fitness coefficients is used to illustrate systematic differences between parameters used to generate data under MutSel and their counterparts estimated by a simple codon model. Motivated by a discrepancy in the literature, interpretation of dN/dS in the context of MutSel is also discussed.

Key words: codon models, population genetics, positive selection, site-specific fitness landscapes, dN/dS .

Introduction

Codon substitution models have provided the basis for the most commonly used methods of inferring positive selection in protein-coding sequences since the pioneering efforts of Muse and Gaut (1994) and Goldman and Yang (1994). Such models produce estimates of the ratio of the nonsynonymous substitution rate (after adjusting for neutral opportunity, dN) to the synonymous substitution rate (likewise adjusted, dS). The rate ratio dN/dS is represented by the parameter ω in a 61×61 substitution rate matrix that is the building block for a variety of popular models (Yang and Nielsen 1998).

The simplest model is M0 (Yang *et al.* 2000), which estimates a single ω for all sites and branches. The limited statistical power of M0 spurred the development of models that account for variation in ω across branches (Yang and Nielsen 1998), across sites (e.g., the M-series models of Yang *et al.* 2000), and across both branches and sites (Guindon *et al.* 2004; Yang *et al.* 2005; Kosakovsky Pond *et al.* 2011; Murrell *et al.* 2015; Smith *et al.* 2015). Positive selection is inferred when a model that permits ω to be >1 fits the data significantly better than a nested version of the same model for which all ω are restrained to be one or less. Such inferences are characteristic of two positive selection scenarios: episodic changes in functional constraints causing transient increases

in ω , and frequency-dependent selection causing sustained elevations in ω along an entire lineage. The signal for episodic selection is typically restricted to a few branches of a phylogeny, and can occur in association with events such as horizontal gene transfer (Yang *et al.* 2013), gene duplication (Pegueroles *et al.* 2013), or colonization of a new niche (Bielawski *et al.* 2004). The signal of frequency-dependent selection differs in that ω is elevated at some sites over much longer periods of evolutionary history. Consequently, frequency dependent selection is easier to detect and has been consistently connected to immune surveillance (Hughes and Nei 1988) and reproductive conflict (Swanson *et al.* 2003), to name two scenarios (see Yang and Bielawski 2000, for a more comprehensive list of examples of this type).

Codon models are phenomenological in the sense that they summarize the “net resultants of selection” (Rodrigue and Philippe 2010) with only limited consideration of generating mechanisms (also see Liberles *et al.* 2013). This approach is necessary when used for inference, largely due to limitations in the information contained in any given alignment. Such models are therefore widely used for the purpose of identifying components of genes that have been impacted by positive selection. They are also used to generate data for the purpose of model testing and comparison. However, models

formulated in terms of the population-level generating mechanisms can also be used to investigate the properties of codon models (Spielman and Wilke 2015). Indeed, without such investigations, we could not understand how phenomenological models are related to the evolutionary processes that actually drive the changes we observe among sequences.

The mutation–selection (MutSel) framework proposed by Halpern and Bruno (1998) provides a mechanistic description of the evolutionary process that is grounded in population genetics. Under this approach, a nucleotide mutation model that is the same for all sites is combined with fixation probabilities computed from site-specific vectors of fitness coefficients assuming a Wright–Fisher population with mutation and selection. Because this framework yields an explicit relationship between dN/dS and fitness differences between amino acids, it provides a principled means to generate realistic sequence alignments consistent with both positive selection scenarios described earlier. Extensions of the MutSel admit the scenario of episodic positive selection by imposing changes in fitness coefficients at specified sites and branches (dos Reis 2015). Frequency-dependent selection can be modeled by causing fitness coefficients to change upon every substitution, as demonstrated in the section titled “M-Series Models Interpreted as Frequency-Dependent Selection under the MutSel Framework” herein. As both scenarios involve temporal changes in fitness coefficients, they can be viewed as cases of population evolution on a changing fitness landscape.

The case where a population evolves on a fixed fitness landscape represents a stark contrast to both the episodic and frequency-dependent scenarios. In the fixed-landscape scenario, the population is not considered to be undergoing adaptation, and an estimate of ω obtained from a phenomenological codon model is expected to be consistent with purifying selection (i.e., $\omega < 1$). Spielman and Wilke (2015) investigated this scenario under the MutSel framework, and asserted that positive selection cannot occur when fitness coefficients are fixed. Indeed, the long-run average dN/dS at a site (and thus ω derived from a codon model) will always be < 1 under MutSel when fitness coefficients are fixed (provided synonymous codons have equal fitness). However, characterizing evolution on a static landscape using the long-run average rate ratio ignores temporal dynamics that arise due to population level processes of mutation and drift.

In this article, we employ the MutSel framework to investigate population level temporal dynamics that can occur on a static fitness landscape, and explore its implications for macroevolutionary inference under phenomenological codon models. We show that a site-specific temporally dynamic dN/dS that depends on the codon occupying the site results from what we call shifting balance (cf. Wright 1982), meaning the process whereby drift causes a site to move away from its fitness peak and a combination of positive selection and drift subsequently act to force it back. We show that the shifting balance phenomenon is detectable by codon models designed to identify site-specific episodic positive selection, and that estimates of ω obtained by these models can be significantly > 1 under certain conditions. This suggests that,

even when fitness coefficients are constant, temporal variations in ω can be misinterpreted as evidence of positive selection due to episodic or frequency-dependent selection, whereby fitness coefficients change over time. Furthermore, we show that the expected proportion of substitutions fixed by positive selection under a shifting balance process (i.e., to compensate for chance substitutions to deleterious amino acids), can be substantial under certain conditions. Taken together, we assert that positive selection can be an important process even when fitness coefficients are constant.

Paper Outline

We start by presenting the elements of the MutSel model following Halpern and Bruno (1998). Motivated by discrepancies in the way dN/dS was formulated by dos Reis (2015) versus Spielman and Wilke (2015), this section includes a discussion as to how the “opportunity” for synonymous and nonsynonymous substitutions can be defined under the MutSel framework. MutSel is then used to validate the interpretation of the M-series models of Yang *et al.* (2000) as being specific for frequency-dependent selection. Although hinted at by other authors (Nielsen and Yang 2003; Kryazhimskiy and Plotkin 2008; Mugal *et al.* 2013), a demonstration of this interpretation has never been published (but see dos Reis 2013, unpublished). We include this because it helps to elucidate the differences between the mechanistic framework and the standard phenomenological approach. We then use the MutSel framework to provide a theoretical explanation for what we call the shifting balance phenomenon, and introduce the notion of a site-specific MutSel landscape. We use two different ways to represent this landscape to illustrate an interesting implication of MutSel that has not been fully appreciated, namely that a site can be occupied by a sub-optimal amino acid for long periods when selection is stringent. This results from what we call a “split” MutSel landscape, the possibility of which has implications for interpreting the signal for functional divergence at the molecular level.

We proceed with the main point of our study by first constructing a mechanistic model for shifting balance. This is used to show that site-specific variations in dN/dS due to shifting balance are most pronounced when the substitution process is not dominated by selection or drift, but admits interplay between the two. We then show that a covarion-like model (CLM3) similar to *fitmodel* (Guindon *et al.* 2004), designed to infer site-specific episodic positive selection, can detect temporal changes in ω generated by shifting balance when this interplay exists. We also show that, under the same conditions, both CLM3 and the branch-site model BUSTED (Murrell *et al.* 2015) can sometimes detect positive selection due to shifting balance. These results suggest that the two models cannot always distinguish between episodic changes in function (i.e., where amino acid fitness have changed) and shifting balance on a static fitness landscape. We also present the MutSel model with peak shifts, following dos Reis (2015), and examine relationships between parameters used to generate data under this model and their counterparts estimated

using the codon model M0. Implications of our results are explored in Discussion.

The Mutation Selection Model

In this section, we first review the mutation–selection framework as presented by Halpern and Bruno (1998). This leads to a novel definition of the site-specific rate ratio. We then justify our novel approach with a discussion concerning the best way to think about substitution rates “before selection” within the MutSel framework.

Let $\mathbf{f}^h = \langle f_1^h, \dots, f_{61}^h \rangle$ be a row vector of “additive” fitness coefficients (Sella and Hirsh 2005) for codons at site h in a n -site protein-coding gene (we assume that mutations to any of the three stop codons are always lethal). Although it is not a requirement, here we assume that the f_j^h are constant across synonymous codons. The probability that a mutation from wild-type codon i to codon j occurs in an individual during the time interval $[0, \Delta t]$ and is subsequently fixed in a haploid population of effective size N_e , assuming that $\mu_{ij}N_e \ll 1$ and $|f_j^h - f_i^h|$ is small, is:

$$\Pr = \begin{cases} \mu_{ij}N_e\Delta t \times \frac{1}{N_e} & \text{if } s_{ij}^h = 0 \\ \mu_{ij}N_e\Delta t \frac{s_{ij}^h/N_e}{1 - \exp(-s_{ij}^h)} & \text{otherwise} \end{cases} \quad (1)$$

where, $s_{ij}^h = 2N_e(f_j^h - f_i^h)$ is the scaled selection coefficient, $\mu_{ij}N_e\Delta t$ approximates the probability that the mutation occurs in one individual during $[0, \Delta t]$, and $(s_{ij}^h/N_e)/(1 - \exp(-s_{ij}^h))$ (or $1/N_e$ when $s_{ij}^h = 0$) approximates the probability that the mutation is eventually fixed (Kimura 1962). We assume that Δt is large enough (e.g., thousands of generations) that the mutant gene is highly likely to be either fixed or eliminated at some time during $[0, \Delta t]$, and that $\mu_{ij}N_e$ is small enough that it is highly unlikely that more than one mutation occurs during the same time interval.

Equation (1) describes a discrete-time Markov process. Probabilities approximate rates on the much larger macroevolutionary time scale, and the elements of the site-specific rate matrix A^h for a continuous-time macroevolutionary substitution process can be defined as follows for all $i \neq j$:

$$A_{ij}^h \propto \begin{cases} \mu_{ij} & \text{if } s_{ij}^h = 0 \\ \mu_{ij} \frac{s_{ij}^h}{1 - \exp(-s_{ij}^h)} & \text{otherwise} \end{cases} \quad (2)$$

The implied proportionality constant scales the rate matrix so that branch lengths are measured as the expected number of single nucleotide substitutions per codon (instead of per generation, see “Methods” section). Diagonal elements A_{ii}^h are specified to make rows sum to zero. It is assumed that $A_{ij} = 0$ unless a single nucleotide change, k to ℓ , changes i to j . The neutral substitution rate is then proportional to μ_{ij} , the rate at which the $k \rightarrow \ell$ mutation occurs under a nucleotide model such as HKY (Hasegawa et al. 1984). It can be shown that A^h defines a time reversible

process so that $\pi_i^h A_{ij}^h = \pi_j^h A_{ji}^h$ for all $i \neq j$ (Halpern and Bruno 1998). Consequently, the row vector $\pi^h = \langle \pi_1^h, \dots, \pi_{61}^h \rangle$ of stationary frequencies for a site is determined by the relation $\pi_i^h \propto \pi_i^0 \exp(2N_e f_i^h)$ where $\pi^0 = \langle \pi_1^0, \dots, \pi_{61}^0 \rangle$ is the vector of stationary frequencies under the neutral model (see supplementary material section, Supplementary Material online).

Let dS^h be the ratio of the synonymous substitution rate (rS^h) to the rate at which synonymous mutations arise at a site (i.e., the “opportunity” for a synonymous substitution, S_{opp}^h). From equation (1):

$$dS^h = \frac{rS^h}{rS_{opp}^h} = \frac{\sum_{(i,j)} \pi_i^h \mu_{ij} N_e \Delta t \times \frac{1}{N_e} I_S}{\sum_{(i,j)} \pi_i^h \mu_{ij} N_e \Delta t I_S} = \frac{1}{N_e} \quad (3)$$

where I_S is an indicator for synonymous (i, j). The ratio of the nonsynonymous substitution rate to the rate at which nonsynonymous mutations arise is similarly expressed:

$$\begin{aligned} dN^h &= \frac{rN^h}{rN_{opp}^h} = \frac{\sum_{(i,j)} \pi_i^h \mu_{ij} N_e \Delta t \times \frac{s_{ij}^h/N_e}{1 - \exp(-s_{ij}^h)} I_N}{\sum_{(i,j)} \pi_i^h \mu_{ij} N_e \Delta t I_N} \\ &= \frac{\sum_{(i,j)} \pi_i^h A_{ij}^h I_N}{\sum_{(i,j)} \pi_i^h \mu_{ij} I_N} \times \frac{1}{N_e} \end{aligned} \quad (4)$$

The site-specific rate ratio is therefore:

$$dN^h/dS^h = \frac{\sum_{(i,j)} \pi_i^h A_{ij}^h I_N}{\sum_{(i,j)} \pi_i^h \mu_{ij} I_N} \quad (5)$$

$dN^h/dS^h = \omega$ when A^h is replaced by the phenomenological rate matrix (e.g., Q of equation (19) in “Methods” section), so equation (5) is the analogue of ω under the mutation–selection framework. Note that, we use ω , and ω_1 without a superscript to refer to a rate ratio shared by a number of sites, whereas superscripts ω^h and ω_1^h indicate site-specific ratios. The dN/dS notation is reserved for theoretical expressions derived under MutSel.

Our definitions for dS^h and dN^h depart from the approach used by both dos Reis (2015) and Spielman and Wilke (2015). To illustrate the differences, it is helpful to first look at how dS has been traditionally defined under a phenomenological codon model. The synonymous substitution rate under the phenomenological rate matrix can be expressed as follows:

$$rS = \sum_{(i,j)} \kappa \pi_i \pi_j I_{TS} + \sum_{(i,j)} \pi_i \pi_j I_{TV} \quad (6)$$

where I_{TS} and I_{TV} are indicators for transitions and transversions, respectively, κ is the transition/transversion ratio, and the π_i are empirically derived codon frequencies. Under this framework, opportunity is defined as the substitution rate between codons that are equally fit, which is obtained by setting $\omega = 1$. But since equation (6) does not contain ω , rS_{opp} is identical to rS and $dS = rS/rS_{opp} = 1$.

Note that rS and rS_{opp} are computed under equation (6) using the same vector of empirical frequencies. Neutral

frequencies under the mutation–selection framework, by contrast, are generally not the same as site-specific frequencies. Perhaps with this in mind, dos Reis (2015) defined dS^h and dS_{opp}^h under the MutSel framework as follows:

$$rS^h = \sum_{(i,j)} \pi_i^h \mu_{ij}^h I_S \quad (7)$$

$$rS_{opp}^h = \sum_{(i,j)} \pi_i^0 \mu_{ij}^h I_S \quad (8)$$

where the π_i^0 are neutral stationary frequencies. By this formulation, rS_{opp}^h can be interpreted as the synonymous substitution rate that would be realized at a site in a pseudo-gene for which all codons have the same fitness (i.e., the substitution rate without selection); rS_{opp}^h is consequently the same for all sites. A different approach was taken by Spielman and Wilke (2015), who defined rS^h as in equation (7), but used site-specific stationary frequencies in place of neutral frequencies in their definition of rS_{opp}^h , making it site-specific. However, this definition negates the interpretation of opportunity as the substitution rate without selection because the π_i^h are determined by selection.

We take a different approach by defining opportunity as the rate at which mutations arise at a site under its current selection regime. It is, after all, the mutations that are actually realized at a site that come under selection. By this definition, rS_{opp}^h is computed using the site-specific stationary frequencies (the π_i^h). Furthermore, rS_{opp}^h now counts mutations, not substitutions. It is therefore derived from the probability of a synonymous mutation ($\propto \mu_{ij} N_e I_S$) instead of the probability of a synonymous substitution ($\propto \mu_{ij} I_S$). By this approach, equation (3) is the ratio of the rate at which synonymous substitutions are realized at a site to the rate at which synonymous mutations arise at the site under the current selection regime specified by a vector of fitness coefficients and N_e . The ratio dS^h is therefore $1/N_e$ instead of one as it is under the traditional phenomenological approach. But note that our approach does not impact the rate ratio dN^h/dS^h , since $1/N_e$ appears as a factor in the expression for dN^h as well. As a result, our equation (5) is the same as that used by Spielman and Wilke (2015); it is only our interpretation of rS_{opp}^h and rN_{opp}^h that is novel.

To reiterate, whereas rN^h/rS^h is the ratio of the expected number of nonsynonymous substitutions to the expected number of synonymous substitutions at a site, rN_{opp}^h/rS_{opp}^h by our definition is the same ratio but for mutations, some of which will be fixed and some lost. Thus, for instance, when the ratio of realized substitutions rN^h/rS^h is less than the ratio of potential substitutions rN_{opp}^h/rS_{opp}^h , the expected proportion of nonsynonymous mutations that are fixed is smaller than the expected proportion of synonymous mutations that are fixed, resulting in $dN^h/dS^h < 1$, a signature of purifying selection. Furthermore, as will be shown in the section titled “MutSel on a Changing Fitness Landscape”, the rate ratio can be transiently >1 following a change in fitness coefficients. Equation (5) is therefore consistent with what the rate ratio was intended to measure.

Results

M-Series Models Interpreted as Frequency-Dependent Selection under the MutSel Framework

The basis of most phenomenological codon models (e.g., the M-series models of Yang et al. 2000) is a 61×61 substitution rate matrix Q with one ω (see “Methods” section). This matrix characterizes the substitution process either for all sites (e.g., as it would under M0) or for some subset of sites (e.g., under M3 sites are apportioned between several ω -categories). Q is similar to the rate matrix for MutSel, as $Q_{ij} = 0$ unless codons i and j differ by a single nucleotide substitution and $Q_{ij} \propto \mu_{ij}$, the neutral substitution rate, when i and j are synonymous. The two rate matrices differ only in their treatment of nonsynonymous substitutions. Rates under Q are scaled by a constant factor ω when i and j are nonsynonymous, so that $Q_{ij} \propto \omega \mu_{ij}$ for all pairs of nonsynonymous codons. By comparison, substitution rates between nonsynonymous codons can be different for each (i, j) pair in the rate matrix A defined by equation (2).

The rate ratio for a set of sites evolving under Q is constantly ω , consistent with the stationary process that Q defines. However, this constancy belies an implicit assumption of a constantly changing fitness landscape. For this reason, $\omega > 1$ under any M-series model can be interpreted as an indication of adaptive evolution by frequency-dependent selection. Here we demonstrate the veracity of this statement using a formal MutSel model for frequency-dependent selection. For an alternative demonstration see dos Reis (2013).

Consider a variation of MutSel where: (i) the incumbent amino acid at a site has one fitness coefficient f^h while all others have fitness $f^h + \Delta f^h$; and (ii) when a substitution occurs, the incumbent and incoming amino acids swap fitnesses so that condition (i) still holds (Nielsen and Yang 2003; Mugal et al. 2013). The vector of site-specific fitness coefficients f^h under assumptions (i) and (ii) is a time-dependent random variable that has no analogue in the ω -model framework. Nevertheless, because the parameters of the process at a site do not change until a substitution occurs, Markov chain properties imply:

- (i) The probability that codon i is substituted by j is proportional to A_{ij}^h , with $A_{ij}^h = \mu_{ij}$ for synonymous and $\omega^h \mu_{ij}$ for nonsynonymous (i, j) , where $\omega^h = 2N_e \Delta f^h / (1 - \exp(-2N_e \Delta f^h))$.
- (ii) For wild-type codon i , the time until a substitution occurs is an exponential random variable with mean $r_i = -1/A_{ii}^h$.

Significantly, properties (i) and (ii) define a Markov process with rate matrix $A^h = Q^h$ (Ross 1996, Chapter 5) (using the formulation with nucleotide frequencies, see “Methods” section). Note that the vector of fitness coefficients f^h at a site is dynamic since it depends on the codon currently occupying the site. It can therefore be different from one site to the next at any instant. All sites that share the same Δf^h nevertheless evolve under the same phenomenological rate matrix Q^h .

The equivalence of the rate matrix A^h for a MutSel process under (i) and (ii) to the rate matrix Q^h (defined in “Methods”

section) suggests that M-series models can be interpreted as being designed to detect signatures of frequency-dependent selection where, for instance, antagonistic interactions between proteins cause the fitness of any given variant to be inversely proportional to its frequency in the population. This interpretation makes sense only when $\Delta f^h > 0$ (i.e., $\omega^h > 1$) however, as was pointed out by dos Reis (2013); it is more appropriate to think of Q^h as a model for purifying selection when $\Delta f^h < 0$ ($\omega^h < 1$), or neutral selection when $\Delta f^h \approx 0$ ($\omega^h \approx 1$). Furthermore, even when $\Delta f^h > 0$, Q^h only captures the phenomenological effect of frequency-dependent selection, the sustained elevation of ω to a value > 1 over a branch or lineage. MutSel under (i) and (ii) is an extension of the two-state model in dos Reis (2013). That model was more mechanistic in the sense that it explicitly accounted for changes in fitness as a function of changes in frequency. Nevertheless, dos Reis (2013) showed that under certain stringent conditions his frequency-dependent model was equivalent to a site evolving under a constant rate ratio. The implication of the interpretation of $\omega > 1$ in an M-series model as an indication of frequency-dependent selection vis-a-vis our study of shifting balance is raised in Discussion.

Shifting Balance on a Static MutSel Landscape

Turning now to the general MutSel model of the section titled “The Mutation Selection Model”, Spielman and Wilke (2015) proved that positive selection signified by a long-run average $dN^h/dS^h > 1$ is not possible when fitness coefficients are fixed (provided that synonymous codons have equal fitness and the mutation process is symmetrical; we establish in supplementary material section, Supplementary Material online, that their result holds for asymmetric mutation processes as well). Here we propose a different interpretation of MutSel that takes into account the temporal dynamics of mutation and drift on a static fitness landscape. The amino acid occupying the h^{th} site will vary over population time scales via mutation–selection as long as at least two amino acids have nonnegligible stationary frequencies. Under this condition, the expected proportion p_+^h of substitutions due to positive selection is:

$$p_+^h = \frac{\sum_{(i,j)} \pi_i^h (A_{ij}^h - \mu_{ij}) I_+}{\sum_{i \neq j} \pi_i^h A_{ij}^h} \quad (9)$$

where I_+ is an indicator for $s_{ij}^h > 0$ (see supplementary material section, Supplementary Material online, for details). This proportion is > 0 unless only one amino acid is viable at the site, which is unlikely to occur at all sites in a gene. The model therefore predicts that some mutations will be fixed by positive selection. Similarly, some mutations will be lost to purifying selection and some fixed by drift. Time reversibility implies that the number of deleterious substitutions is exactly balanced by an equal number that are beneficial when averaged over population time scales (e.g., $p_+^h = p_-^h$, where p_-^h is the proportion of substitutions $i \rightarrow j$ for which $s_{ij}^h < 0$). Hence, beneficial substitutions can be thought of as “repairing” previously deleterious ones (Sella and Hirsh 2005; Mustonen and Lässig 2009).

To visualize this process, we introduce the notion of a site-specific MutSel landscape (cf. Bazykin 2015), an analogue of the traditional fitness landscape constructed by sorting site-specific stationary frequencies from largest to smallest as depicted in figure 1. There, the frequencies were derived from fitness coefficients drawn from a normal distribution with a standard deviation $\sigma = 0.001$ (we use an effective population size of $N_e = 1,000$ in all of our analyses unless otherwise indicated). Positive selection can be seen to occur on this landscape by considering how dN^h/dS^h varies over time. The dN^h/dS^h ratio for the site depicted in figure 1 is 0.58. However, this value is a long-run average; the rate ratio in fact varies depending on the codon currently occupying the site. The codon-specific dN_i^h/dS_i^h (where i is the codon currently occupying the site) can be computed from the i^{th} row of the rate matrix:

$$dN_i^h/dS_i^h = \frac{\sum_j A_{ij}^h I_N}{\sum_j \mu_{ij} I_N} \quad (10)$$

Equation (10) is approximately equal to ω when A^h is replaced by the rate matrix for M0, with minor variations due to biases in the mutation process (e.g., the proportion of single-nucleotide nonsynonymous substitutions from i to j that are transitions varies slightly between codon pairs when $\kappa > 1$). However, dN_i^h/dS_i^h can undergo substantial variations under MutSel, depending on the fitness coefficients.

The line plot in figure 1 shows dN_i^h/dS_i^h (scaled on the right y-axis) for each codon. When a codon with low fitness (one far to the right or in the “tail” of the MutSel landscape) occupies the site, the majority of mutations are “up-slope” with $s_{ij}^h > 0$. The codon-specific rate ratio (eq. 10) is consequently greater than $dN^h/dS^h = 0.58$ (as large as 4.10 in this example). As the site moves in the up-slope direction, the proportion of mutations that are further up-slope diminishes and dN_i^h/dS_i^h decreases to a value < 0.58 (as small as 0.24 in this example). By this process, chance substitutions (e.g., drift) that move a site down-slope are balanced by a combination of drift and positive selection ($dN_i^h/dS_i^h > 1$) that move the site back toward its peak. We suggest calling temporal variation in the site-specific rate ratio




FIG. 1. Bars show stationary frequencies sorted from largest to smallest. The line shows the codon-specific rate ratio dN_i^h/dS_i^h for the sorted codons. The rate ratio varies depending on the codon currently occupying the site, and can be > 1 following a chance substitution into the tail (to the right) of the landscape. In this case the codon specific rate ratio for the site ranges from 0.24 to 4.10 with a temporal average of $dN^h/dS^h = 0.58$.

“shifting balance” because the process is evocative of Wright’s theory (Wright 1982).

Split MutSel Landscapes

Sites are considered to be homogeneous under a codon substitution model such as M0. MutSel, by contrast, is site-specific. Here we make use of this property to investigate the dynamics of the MutSel substitution process at an individual site. We show that a site-specific landscape can be quite complex, and that the temporal dynamics of evolution implied by such a landscape can change dramatically with changes in population size alone (i.e., even while fitness coefficients are fixed).

Under a stringent selection regime, a landscape can be very sparse; that is, it is possible for the site-specific vector of stationary frequencies to be nearly zero for all but a few amino acids. Some sparse landscapes can be structured such that the shifting balance process can reside for nonnegligible periods of time at sub-optimal amino acids. We call these “split landscapes” because a sequence alignment derived from evolution on such a landscape will tend to consist of site patterns that are split between the optimal and a sub-optimal amino acid.

The method of visualizing a landscape introduced by McCandlish (2011) is particularly useful in this context (see “Methods” section). An example is shown in figure 2. There, only three amino acids occur with nonnegligible frequencies (fig. 2A), all of which are in the region we call the peak of the landscape, where $dN_i^h/dS_i^h < 1$ for all codons i . The vertices in figure 2B mark locations of the eight codons for these amino acids in the 2-dimensional landscape (codons with negligible stationary frequencies were omitted for clarity). The diameter of each circle is proportional to the stationary frequency of the codon at its center. The length of each line connecting a pair of vertices is approximately proportional to

the expected number of single nucleotide substitutions it would take for the site to move from one vertex to the other and back again via any and all available pathways. Pairs of codons that are very close to one another (e.g., ACA and ACG) are synonymous and differ by one transition; they are close because the site can move rapidly between them.

The depicted site will spend most of its time moving between T (with stationary frequency $\pi_T = 0.76$), E ($\pi_E = 0.19$), and K ($\pi_K = 0.05$), since all other amino acids have negligible frequencies. The dynamic is largely governed by that fact that T and E differ by at least two nucleotides. A site that starts at one must therefore be occupied by an alias of K before reaching the other. However, K is the least fit of the three, and so can only be reached after a long expected waiting time. This means that the site can become stuck at the sub-optimal E for extended periods. This effectively “splits” the MutSel landscape. Note that allowing double and triple mutations would not necessarily change this dynamic but would only replace low substitution rates between T and E via K with low rates of double and triple mutations that take a site from T (or E) directly to E (or T) (for example, Tamuri et al. (2012) estimated that double mutations comprise 0.58%, and triple only 0.002% of all mutations in their analysis of 12 proteins in the mitochondrial genome of 244 mammal species). The landscape remains split either way.

To highlight the role of drift on a static landscape, we reduced the effective population size N_e by a factor of 10, which allows drift a much larger role in the evolutionary process. Figure 3 depicts the landscape after this reduction (note that the fitness coefficients are the same as those used in fig. 2). The MutSel landscape shown in figure 3A is now much broader, reflecting an increase in frequencies that were previously negligible. The 2-dimensional landscape in figure




Fig. 2. A MutSel landscape for a site under stringent selection pressure. A: Only three amino acids have nonnegligible frequencies. B: The 2-dimensional landscape provides information about the substitution dynamics at the site. Vertices indicate the location of codons in the 2-dimensional space; circle diameters are proportional to stationary frequencies; edge lengths are proportional to the expected number of single nucleotide substitutions required for the site to move from one vertex to the other and back again.

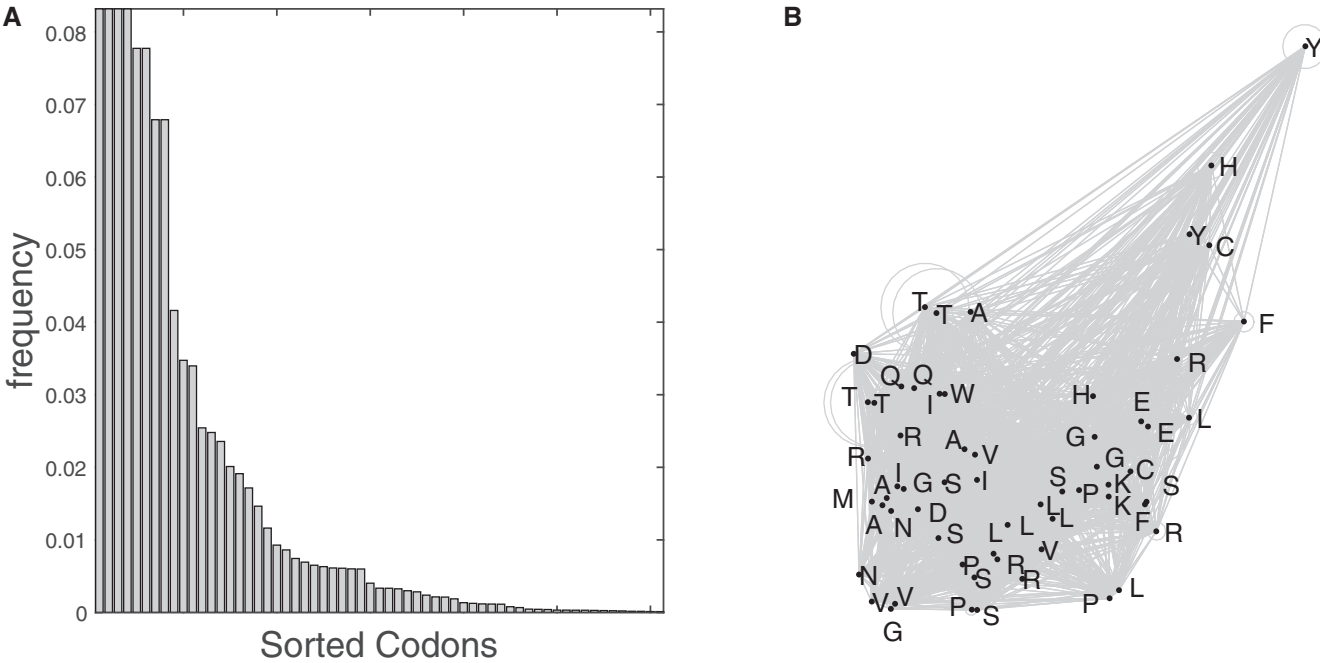


Fig. 3. The landscape in figure 2 after a 10-fold decrease in the effective population size. A: The MutSel landscape is now broader; many codons have nonnegligible stationary frequencies. B: The site is now free to move around a large network of connections.

3B depicts a much larger network of connections between viable amino acids. By increasing the role of drift, we reduced the effectiveness of selection, and thereby freed the site to move more rapidly between *T* and *E* (third column, table 1), whose stationary frequencies are now nearly the same ($\pi_T = 0.083$, $\pi_E = 0.078$). This example demonstrates that site-specific codon frequencies and the expected rate of nonsynonymous substitutions (table 1) can change dramatically over time due to changes in effective population size while the underlying relationship of the site to protein function (i.e., the site’s fitness coefficients) remains constant.

A Mechanistic Model for Shifting Balance

It is reasonable to expect branch-site codon models (Guindon *et al.* 2004; Yang *et al.* 2005; Kosakovsky Pond *et al.* 2011; Murrell *et al.* 2015; Smith *et al.* 2015) to capture temporal variations in site-specific rate ratios illustrated in the section titled “Shifting Balance on a Static MutSel Landscape” when fitted to data generated under MutSel. In this section, we

derive expressions from the mutation–selection framework for parameters that have meaningful interpretations in the context of a branch-site model that allows a site to switch between two rate ratios, ω_1 and ω_2 , continuously over time at a rate δ (Guindon *et al.* 2004). The purpose of this exercise is 2-fold: first, to demonstrate that there is a mechanism by which a site can switch between two rate ratios on a static landscape; and second, to identify the conditions under which such variations are most pronounced.

Let I_p^h be an indicator for codons *i* for which $dN_i^h/dS_i^h \leq 1$ (e.g., near the peak of the MutSel landscape), and let I_t^h be the same for codons for which $dN_i^h/dS_i^h > 1$ (in the landscape’s tail). A site will shift between its peak and tail, corresponding to switches between $\omega_1^h \leq 1$ and $\omega_2^h > 1$, with stationary proportions

$$p_1^h = \sum_i \pi_i^h I_p^h \tag{11}$$

$$p_2^h = 1 - p_1^h \tag{12}$$

Expected rate ratios can be computed using equation (5) restricted to either the peak or tail:

$$\omega_1^h = \frac{\sum_{(i,j)} \frac{\pi_i^h}{p_1} A_{ij}^h I_N^h I_p^h}{\sum_{(i,j)} \frac{\pi_i^h}{p_1} \mu_{ij} I_N^h I_p^h} \tag{13}$$

$$\omega_2^h = \frac{\sum_{(i,j)} \frac{\pi_i^h}{p_2} A_{ij}^h I_N^h I_t^h}{\sum_{(i,j)} \frac{\pi_i^h}{p_2} \mu_{ij} I_N^h I_t^h} \tag{14}$$

NOTE.—The Codon-specific rate ratios for each codon listed in the first column as a function of the effective population size N_e . The site-specific MutSel landscapes for $N_e = 1,000$ and $N_e = 100$ are shown in figures 2 and 3.

The expected number of switches per single nucleotide substitution is the ratio:




FIG. 4. The distributions of the site-specific rate ratio dN^h/dS^h for 250 sites under each of the three values of σ used in this study. A: $\sigma = 0.0001$ is consistent with nearly neutral evolution, as most sites evolve under a rate ratio close to one. B: the rate ratio varies over a broad range of values when $\sigma = 0.001$. C: $\sigma = 0.01$ is consistent with stringent selection for which most sites evolve under a rate ratio close to zero.

$$\delta^h = \frac{\sum_{(i,j)} \pi_i^h A_{ij}^h I_{\text{switch}}}{\sum_{i \neq j} \pi_i^h A_{ij}^h} \quad (15)$$

where I_{switch} is an indicator for pairs of codons (i, j) for which one is in the peak and the other the tail of the MutSel landscape. Since a switch can only occur upon a substitution, δ^h can be no greater than one.

Fitness coefficient vectors were drawn for 250 sites from a multivariate normal distribution centered at zero and with covariance matrix $\sigma^2 I$, where I is the 61×61 identity matrix, using $\sigma \in \{0.0001, 0.001, 0.01\}$ and $N_e = 1,000$; the mutation process was assumed to be HKY with $\kappa = 4$ and uniform nucleotide frequencies (see “Methods” section). Each vector was modified to make synonymous codons have the same fitness. The site-specific rate ratio was then calculated for each vector. Figure 4 shows how the distribution of dN^h/dS^h changes with σ . Figure 4A demonstrates that $\sigma = 0.0001$ corresponds to a nearly neutral selection regime, as the site-specific rate ratio is very nearly one most of the time (with a median dN^h/dS^h of 0.99). Figure 4B shows that sites evolve over a wide range of site-specific rate ratios when $\sigma = 0.001$ (dN^h/dS^h ranges from 0.06 to 0.74 with a median of 0.39). And sites are mostly under stringent selection in figure 4C, when $\sigma = 0.01$ (with a median dN^h/dS^h of 1×10^{-5}). In the remainder of this section, we characterize the differences between these three selection regimes in terms of ω_1^h , ω_2^h , p_2^h , and δ^h .

Figure 5 shows box plots for parameter values computed from an additional draw of 250 vectors of fitness coefficients for each value of σ . We begin by considering the case of nearly neutral evolution ($\sigma = 0.0001$). The median expected proportion of single nucleotide substitutions attributed to positive selection (p_+^h , fig. 5A) is only 3.4%. The median probability that a site is in the tail of its MutSel landscape (p_2^h , fig. 5B) is just under 43%. The median switching rate (δ^h , fig. 5C) is 0.44, indicating approximately one switch for every two substitutions. The rate ratio from the tail (ω_2^h , fig. 5D) is tightly distributed around a median of 1.1; ω_1^h (not shown) is

similarly distributed but with a median of 0.91. In this scenario the population moves easily in and out of the tail, with only a small difference between ω_1^h and ω_2^h , because the landscape is nearly flat, with only slight variation in the π_i^h and dN_i^h/dS_i^h across codons. Although transient changes in ω^h are expected to occur frequently under this process, they would be deemed trivial in magnitude within a real dataset.

Next, we consider the case where a population is being held tightly to its fitness peak ($\sigma = 0.01$). The median value p_2^h is much $< 1\%$, reflecting a low probability of drift away from the peak. The rate ratio ω_2^h from the tail is relatively large, with median 2.7 (the median of ω_1^h was $< 1 \times 10^{-4}$). The median proportion of single nucleotide substitutions due to positive selection p_+^h and the median switching rate δ^h are both very low. This scenario is consistent with strong selective pressure that mostly prevents the shifting balance process from moving the population away from its peak.

Interestingly, some parameters have outliers under this scenario. Outliers in ω_2^h tend to correspond to cases where p_2^h is very small. Among the 53 trials for which ω_2^h was > 5 , for example, the median value of p_2^h was $< 2 \times 10^{-10}$. Such values indicate very strong selection pressure following a very rare chance substitution into the tail. Outliers in δ^h and p_+^h can be attributed to chance relationships in a reduced space of viable codons (i.e., in a sparse landscape). An example is depicted in figure 6, where codons were sorted by dN_i^h/dS_i^h rather than frequency so that the point of separation between peak and tail could be represented (vertical dashed line). Two amino acids M and I dominate the landscape; almost all substitutions are between them and consist of single nucleotide substitutions in the third position. Since M has only one codon alias, any substitution from the peak can only be nonsynonymous. And since a substitution from M across the $dN_i^h/dS_i^h = 1$ boundary to ATA(I) is quite likely (e.g., table 2), δ^h and p_+^h are both unusually large for the $\sigma = 0.01$ scenario ($\delta^h = 0.58$ and $p_+^h = 0.28$).

Lastly, we consider the case between the nearly neutral and stringent selection scenarios where $\sigma = 0.001$. Population dynamics on this landscape lead to a relatively large median

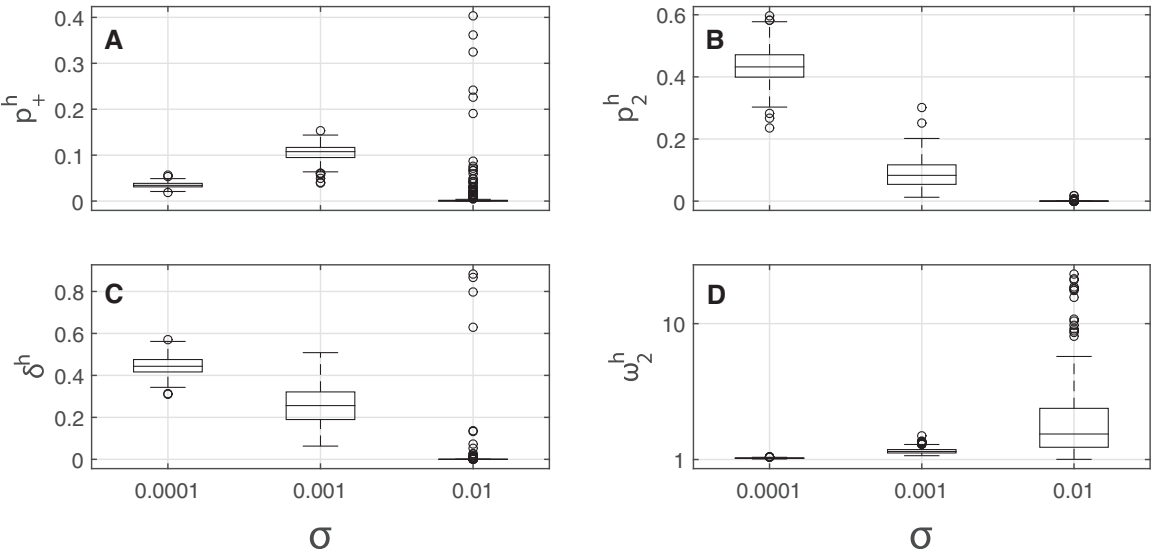


FIG. 5. Distributions for the parameters of the mechanistic shifting balance model computed from 250 trials. Each set of parameters was derived from a vector of fitness coefficients drawn from a normal distribution with zero mean and standard deviation $\sigma \in \{0.0001, 0.001, 0.01\}$ and with $N_e = 1000$. The center bar in each box indicates the median. The limits of each box show the first and third quartiles. Circles indicate outliers. A: p_1^h , the expected proportion of substitutions due to positive selection; B: p_2^h , the expected proportion of time (single nucleotide substitutions) a site is in the tail of its MutSel landscape; C: δ^h , the expected rate a site switches between the peak and tail of its MutSel landscape; D: ω_2^h , the expected rate ratio for the site when it is in the tail of its MutSel landscape.




FIG. 6. The MuSel landscape depicted here was constructed by sorting codons by dN_i/dS_i . The vertical dashed line shows the point of separation between peak and tail. The site is dominated by substitutions between M and the three codons for I. Although substitutions from M are rare, when they do occur they are nonsynonymous since M has only one alias, and they are almost always to I. Substitutions to ATA(I), to the right in the tail of the landscape, are favored due to transition bias. As a result, both $\delta^h = 0.58$ and $p_+^h = 0.28$ are unusually large for this scenario.

Table 2. Substitution Probabilities.

	ATG(M)	ATT(I)	ATC(I)	ATA(I)
ATG(M)	0	0.30	0.30	0.39
ATT(I)	0.76	0	0.17	0.07
ATC(I)	0.76	0.17	0	0.07
ATA(I)	0.87	0.06	0.06	0

NOTE.—Numbers give the probabilities that the incumbent codon in a row is next substituted by the codon in a column for the MutSel landscape depicted in figure 6.

value of p_+^h , which indicates that $\sim 10\%$ of single nucleotide substitutions are due to positive selection acting to “repair” deleterious substitutions. The median switching rate is 0.26, one switch for every four substitutions. The median rate ratio

from the tail is 1.5 (its 0.33 from the peak). Whereas the previous two scenarios represent extreme cases where one population genetic process strongly dominated the other (e.g., drift dominated in the $\sigma = 0.0001$ case, and selection dominated in the $\sigma = 0.01$ case), this scenario reflects an interplay between both processes. Here, the population occasionally moves away from its peak, but such events are quickly corrected because selection remains an effective force for moving the population back. This is the scenario that produces the strongest transient signature of positive selection on a fixed landscape.

The way fitness coefficients are selected introduces a phenomenological component to the mutation–selection framework, since we use an assumed distribution in lieu of actual values. Although a few investigations suggest that the distribution of the s_{ij}^h (and therefore the f_i^h) is sometimes consistent with a normal distribution (Nielsen and Yang 2003; Tamuri et al. 2012), there is no reason not to try alternatives. When we reproduced figure 5 using an exponential distribution with variance $\mu^2 \in \{1 \times 10^{-8}, 1 \times 10^{-6}, 1 \times 10^{-4}\}$, we found similar patterns as those described earlier (supplementary fig. S1, Supplementary Material online). We expect that, whatever distribution is used to draw fitness coefficients, a lower variance corresponds to the nearly neutral scenario dominated by drift, a higher variance to the stringent scenario where selection dominates, and something in between to a balance between selection and drift for which shifting balance is strongest.

Detecting Transient Changes in ω Caused by Shifting Balance on a Fixed Landscape

In the previous section, we investigated a mechanistic process by which a site can switch between two rate ratios as it moves

over its fixed MutSel landscape. Our objective in this section is to demonstrate that a version of the switching model of Guindon *et al.* (2004), suitably modified to resemble the mechanistic switching model of the previous section, can detect site-specific variations in the rate ratio under certain conditions. We call our phenomenological switching model CLM3 for covarion-like M3 (see “Methods” section for details).

Alignments were generated on an 8-taxa symmetrical tree with branch lengths $b \in \{0.25, 0.5, 1\}$ using fitness coefficients with $\sigma \in \{0.0001, 0.001, 0.01\}$ and $N_e = 1,000$ as described in “Methods” section. For each scenario defined by (σ, b) the same set of 500 fitness coefficient vectors was used to generate 50 unique alignments. The data were fitted to CLM3 to obtain maximum likelihood estimates (MLEs) of model parameters. Each alignment was also fitted to codon model M3 (see “Methods” section) to provide a test for the significance of site-specific switches between ω_1 and ω_2 . Table 3 shows the number of trials out of 50 for which the M3-CLM3 contrast rejected the null hypothesis of no switching. The test, conducted at the 5% level of significance, seldom detected evidence for switching under the nearly neutral ($\sigma = 0.0001$) and stringent selection ($\sigma = 0.01$) scenarios, the exception being the $b = 1$ with $\sigma = 0.01$ scenario, where the test was significant in 15 of the 50 trials. Shifting was detected in all trials when $\sigma = 0.001$ and $b \in \{0.5, 1\}$, and in most trials when $b = 0.25$. These results are in agreement with the mechanistic model that indicated that the scenario where neither drift nor selection dominate ($\sigma = 0.001$) would produce the strongest signal for shifting balance.

Previous investigations indicated that a covarion-like model can detect switching (i.e., $\delta > 0$) even when data is generated without switching, and that this may occur when the number of ω -categories used to generate the data is greater than the number assumed by the fitted model (Lu and Guindon 2013). Under our generating scenario with $\sigma = 0.001$, the site specific rate ratio can vary greatly, for example with values as small as $dN^h/dS^h = 0.06$ and as large as $dN^h/dS^h = 0.74$ for the 250 trials depicted in figure 4B. To rule out the possibility that this variation produced false signatures of switching, an additional set of fitness coefficient vectors was drawn with $\sigma = 0.001$. The rate ratio dN^h/dS^h was computed for each vector. Each rate ratio was used to construct a site-specific phenomenological substitution rate matrix (e.g., Q defined in “Methods” section). The resulting generating model was thus similar to an M-series model but

with a different ω^h for each site. This model was used to generate fifty 500-codon alignments on a symmetrical 8-taxa rooted tree with all branch lengths $b = 1$. Since each site was evolved under its own rate matrix with $\omega^h = dN^h/dS^h$, the alignments had a similar distribution of site-specific rate ratios as data generated under MutSel but without the rate shifts that can occur under MutSel. The M3-CLM3 contrast failed to reject the null at the 5% level of significance in all 50 trials, indicating no detectable switching (i.e., δ was never significantly > 0). We concluded that the results in table 3 can be attributed to rate shifts caused by shifting balance.

We used the $\sigma = 0.001$ and $b = 1$ scenario to investigate relationships between MLEs under CLM3 and site-specific parameters ω_1^h , ω_2^h , p_2^h , and δ^h of the mechanistic switching model. Box plots in figure 7 depict the distributions for these parameters computed from site-specific fitness coefficients. A different draw of fitness coefficient vectors was used to generate an 8-taxa alignment for each of the five trials depicted. Diamonds mark MLEs obtained by fitting each alignment to CLM3. Whereas the box plots indicate similar distributions for the site-specific parameters in each trial, the MLEs exhibit substantial variation. They also appear to be correlated such that an alignment that produced a larger estimate for p_2 also produced smaller estimates for both ω_2 and δ . This pattern is consistent with correlations between the 2,500 values computed under the mechanistic switching model (5 alignments, each with 500 sites), for which $r(p_2^h, \omega_2^h) = -0.66$, $r(p_2^h, \delta^h) = -0.52$, $r(\omega_2^h, \delta^h) = 0.84$. The correlations between the MLEs for the 50 trials of the $\sigma = 0.001$, $b = 1.0$ scenario in table 3 were $r(p_2, \omega_2) = -0.86$, $r(p_2, \delta) = -0.94$, $r(\omega_2, \delta) = 0.87$ (supplementary fig. S2, Supplementary Material online).

It is interesting that the MLE for ω_1 is consistently underestimated compared with the median value of ω_1^h , as shown in figure 7A. Here we attempt to explain this result. First note that, even over a branch as long as one single nucleotide substitution per codon, a substantial proportion of sites will remain unchanged. In a simulation of a 500-codon sequence evolving under the $\sigma = 0.001$ scenario over a branch length of one, for example, we found that no substitutions occurred at $\sim 40\%$ of sites (supplementary fig. S3, Supplementary Material online). It is possible that CLM3 underestimates ω_1 relative to the median of ω_1^h in part to account for such sites.

Whereas the five trials in figure 7 each had their own draw of site-specific fitness coefficients, the same set of 500 fitness coefficient vectors (but different than any of those in fig. 7) generated all 50 8-taxa alignments used to produce the box plots in figure 8. There, the box plots on the left in each panel show the distributions for ω_1^h , ω_2^h , p_2^h , and δ^h for the 500 vectors. Box plots on the right show the distributions for the 50 MLEs for each parameter estimated using CLM3. The larger variance exhibited by the MLEs compared with their corresponding site-specific values is likely due to random differences between the generated alignments. As in figure 7A, the MLE for ω_1 is substantially underestimated compared with the median value of ω_1^h . The distribution for the MLEs for ω_2

Table 3. Results, M3-CLM3 Contrast.

b/σ	0.0001	0.001	0.01
1.00	0 (0.51,1.2)	50 (0.00,0.77)	15 (0.00,0.06)
0.50	1 (0.75,1.2)	50 (0.12,0.68)	1 (0.00,0.08)
0.25	2 (0.69,1.1)	39 (0.05,0.79)	1 (0.00,0.08)

NOTE.—The left-most column gives the branch length and the top-most row the value of σ used to generate 50 alignments for the nine (b, σ) scenarios. Each cell shows the number of cases out of 50 for which the M3-CLM3 contrast detected site-specific switches between ω_1 and ω_2 . The numbers inside the brackets are the median values of the MLEs for ω_1 and ω_2 .

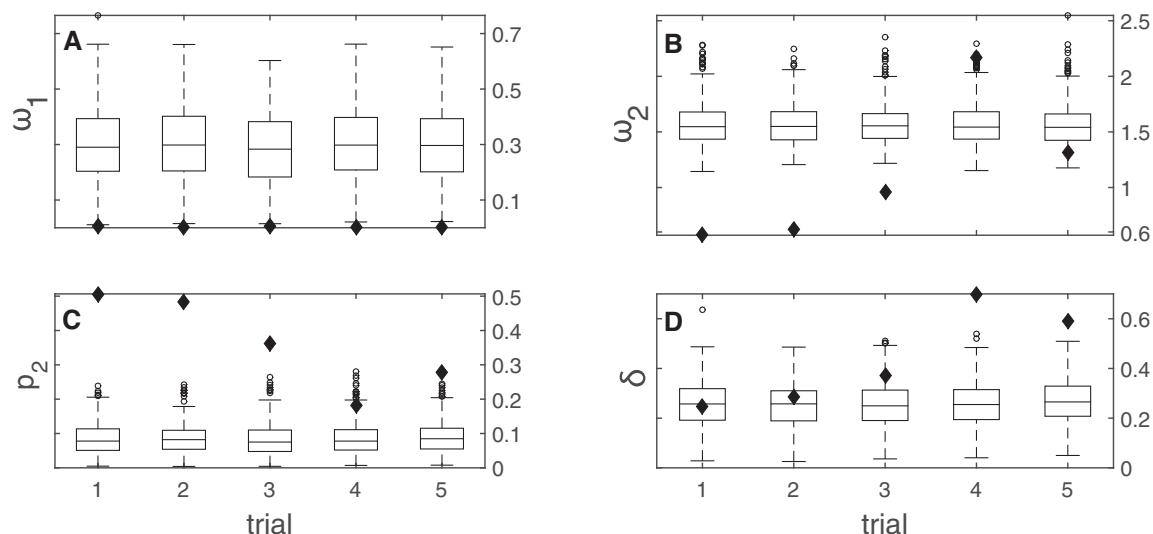


FIG. 7. Diamonds mark maximum likelihood estimates for A: ω_1 , B: ω_2 , C: p_2 , and D: δ , obtained by fitting each of five 8-taxa alignments to CLM3. Alignments were generated under MutSel with $\sigma = 0.001$ and $b = 1$; a different set of 500 fitness coefficient vectors was used for each of the five trials. Box plots show distributions for corresponding site-specific parameters ω_1^h , ω_2^h , p_2^h , and δ^h computed using the mechanistic switching model.

is centered near the median value of ω_2^h . The MLEs for p_2 and δ appear to be systematically larger than the median of the p_2^h and δ^h .

The bias in the median MLEs for ω_1 , p_2 and δ compared with the medians of ω_1^h , p_2^h and δ^h shown in figure 8 can be attributed in part to differences between the MutSel generating process and the assumptions made under CLM3. In particular, sites each have their own parameters under MutSel, whereas they are assumed to share parameters under CLM3. Furthermore, CLM3 assumes a common vector of stationary codon frequencies, whereas sites under MutSel have different frequencies. It is not clear how to account for these differences between the generating process and the fitted model. It is also possible that some of the discrepancy can be attributed to the small number of taxa used in

our study. It is nevertheless interesting that the phenomenological estimate of δ (as an average over sites) was sometimes close to the median value computed using equation (15). Taken together with the results of the M3-CLM3 contrasts in table 3, these findings support the notion that shifting balance on a fixed landscape can lead to a phenomenological signal that is detectable as site-specific changes in ω , and that this is most likely to occur when drift and selection are both at play, as is the case under the $\sigma = 0.001$ scenario.

Detecting Positive Selection Caused by Shifting Balance

In the last section, we focused on the ability of CLM3 to detect transient changes in ω due to shifting balance. Here we shift our attention to the possibility that phenomenological codon

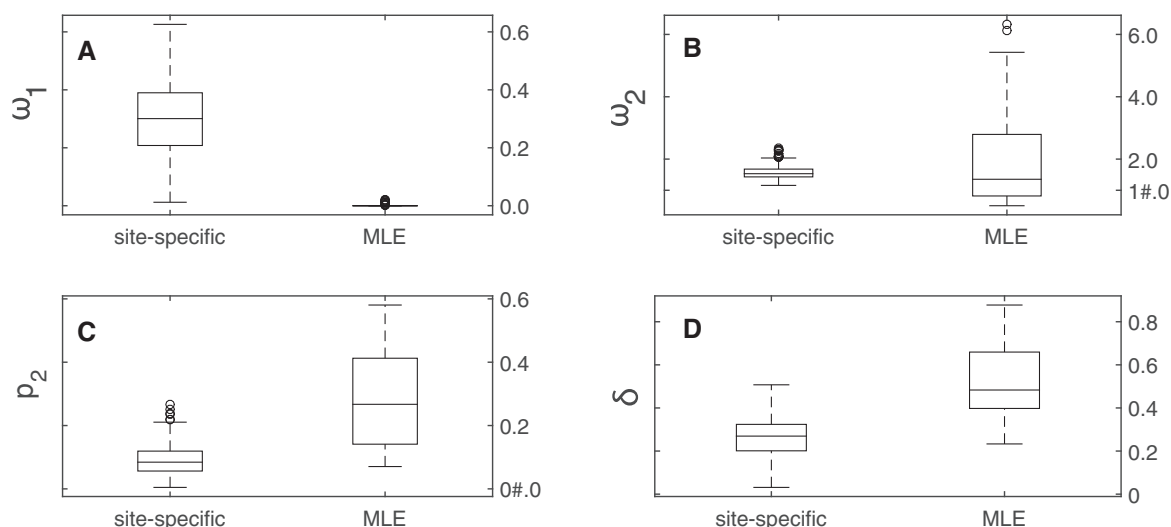


FIG. 8. Fifty 8-taxa alignments were generated under MutSel with $\sigma = 0.001$ and $b = 1$ using the same set of 500 vectors of site-specific fitness coefficients. Panels compare the distribution of site-specific parameters with the distribution of the MLEs for A: ω_1 , B: ω_2 , C: p_2 , and D: δ . MLEs were obtained by fitting the alignments to CLM3.

models might detect transient signals for $\omega > 1$ under this process. In addition to a new test for positive selection based on CLM3, we extend our investigation to include a popular analytical framework called BUSTED (Murrell *et al.* 2015).

First we must place restrictions on the parameters of CLM3 so that it can be used as the basis of a test for switches to $\omega > 1$. In the basic form CLM3 allows sites to switch between two ω -categories at a rate of δ switches per single nucleotide substitution. In the null model (CLM3a) for the new test, CLM3 is restricted such that $\omega_1 < 1$ and $\omega_2 = 1$; thus, positive selection is not permitted. In the alternative model (CLM3b), the value of ω_2 is constrained so that it must be > 1 . As CLM3a and CLM3b are nested, this model pair comprises a likelihood ratio test for episodic positive selection. BUSTED (Murrell *et al.* 2015) also allows the rate ratio at a site to change over time. This model estimates a distribution for ω , consisting of three rate ratios ω_0 , ω_1 , and ω_2 and two proportions p_1 and p_2 , that is shared by all sites. Sites switch between rate-ratios from branch-to-branch. That is, unlike CLM3 where a site can switch its ω -category multiple times along a branch, the rate ratio for a site is constant along any given branch under BUSTED. The null hypothesis under BUSTED is that $\omega_0 \leq \omega_1 \leq \omega_2 = 1$. This is contrasted with the alternative that $\omega_0 \leq \omega_1 \leq 1 \leq \omega_2$.

The same 50 alignments for each of the nine (σ, b) scenarios that were used in the previous section were used here. The test for positive selection under both CLM3 and BUSTED was conducted only if the maximum likelihood estimate for ω_2 was > 1 when ω_2 was unrestricted. All tests were conducted at the 5% level of significance. Table 4 shows the number of trials in each scenario for which BUSTED and the CLM3a-CLM3b contrast found evidence of positive selection. Both models inferred positive selection in substantially $> 5\%$ of the trials, as would have been expected under the null, when $\sigma = 0.001$ and $b = 1$. BUSTED also detected positive selection in the same number of cases under the same σ with $b = 0.5$. Both models detected positive selection in data generated under the $\sigma = 0.0001$ and $\sigma = 0.01$ scenario at a rate consistent with what would be expected by chance under each of their null models (i.e., close to 5% or 2 to 3 trials out of 50), except that BUSTED found signal in 20% of trials in the $\sigma = 0.01, b = 1$ scenario. The distribution of the MLEs for ω_2 and p_2 among the trials for which the null was rejected are shown in figure 9, with the exception of 13 trials for which the MLE for ω_2 under BUSTED was > 50 (supplementary table S1, Supplementary Material online). These results demonstrate that both models can detect the phenomenological signature of positive selection due to shifting balance under the scenario where neither selection nor drift dominate. Subsequent to our analysis, we also fitted the 50 trials of each of the nine (σ, b) scenarios to aBSREL (Smith *et al.* 2015), an alternative branch-site model similar to BUSTED but with fewer constraints. Results show that aBSREL can also detect positive selection due to shifting balance, especially under the $\sigma = 0.001, b = 1$ scenario—see supplementary table S2, Supplementary Material online).

Table 4. Results, CLM3a-CLM3b and BUSTED.

b/σ	0.0001	0.001	0.01
1.00	(1, 3)	(20, 11)	(10, 3)
0.50	(1, 2)	(20, 1)	(3, 1)
0.25	(2, 3)	(5, 0)	(0, 0)

NOTE.—The left-most column gives the branch length and the top-most row the value of σ used to generate 50 alignments for the nine (b, σ) scenarios. Each cell shows the number of cases (x, y) out of 50 for which positive selection was detected by BUSTED (x) and the CLM3a-CLM3b (y) contrast.




Fig. 9. Box plots show the distribution for the MLEs of A: ω_2 and B: p_2 estimated under BUSTED and CLM3b for the trials where the null hypothesis of no positive selection was rejected under each model. The plots for BUSTED do not show 13 trials for which ω_2 was > 50 .

We used branch lengths $b \geq 0.25$ single nucleotide substitutions per codon in all of our trials. Positive selection by shifting balance may be less evident when branches are shorter. In an additional trial, for example, we generated 50 alignments on a 16-taxa tree with $b = 0.10$ on all branches. The M3-CLM3 contrast detected significant shifting in 39 of the trials. However, although the MLE for ω_2 was > 1 in 10 of those trials, it was never significantly > 1 when the CLM3a-CLM3b contrast was applied. BUSTED also did not infer positive selection in any of the alignments. This suggests that the potential of incorrectly attributing positive selection by shifting balance to a change in fitness coefficients is most pronounced when branch lengths are at least ~ 0.25 single nucleotide substitutions per codon.

MutSel on a Changing Fitness Landscape

Up to this point we have shown that the substitution process at a site under the mutation–selection framework can be dynamic or, to put it another way, transiently nonstationary, even if site-specific landscapes are fixed. By contrast, an episodic shift in fitness landscapes can produce a long-lasting nonstationary response if such shifts occur at a number of sites at the same time. The dynamic following a change in fitness landscape was recently illustrated by dos Reis (2015) (also see Mustonen and Lässig 2009). Under his environmental shift (MutSelES) model, A^h is the rate matrix defining the selection regime for the h^{th} site of an ancestral sequence. At

$t = 0$ the regime switches to a different matrix B^h . This change initiates a nonstationary substitution process characterized by an elevated rate ratio $dN^h(t)/dS^h(t)$ that can be quantified in the following way:

$$dN^h(t)/dS^h(t) = \frac{\sum_{(i,j)} \pi_i^h(t) B_{ij}^h I_N}{\sum_{(i,j)} \pi_i^h(t) \mu_{ij} I_N} \quad (16)$$

where $\pi^h(0)$ is the row vector of stationary frequencies for the site consistent with A^h and $\pi^h(t) = \pi^h(0) \exp(B^h t)$ described the process of convergence to the stationary frequencies for B^h . The expected rate for the site at $t = 0$ can be >1 , but will decay exponentially over time until it reaches a value <1 consistent with shifting balance on its new landscape.

Here we consider what happens when M0 is fitted to data generated by the nonstationary process that follows simultaneous changes in fitness coefficients at all sites. Modeling a nonstationary process as stationary can be thought of as a way of estimating an average effect, so it is reasonable to interpret estimates of ω under M0 as a mean taken across all codon sites (n) and the branch length (b). One possible way to formulate this under MutSelES is:

$$\bar{\omega}(b) = \frac{1}{b} \int_0^b \frac{1}{n} \sum_{h=1}^n dN^h(t)/dS^h(t) dt \quad (17)$$

This provides a way to predict the estimate for ω under M0.

To compare predictions with MLEs, 200 sequence pairs (S_1 , S_2) 1,000 codons in length were generated under MutSelES with $\sigma = 0.001$ and $N_e = 1,000$, and with branch lengths ranging between 0 and 1. Each site had its own pair of rate matrices A^h and B^h . Figure 10A shows that the median MLE for ω estimated by M0 is highly correlated with the prediction $\bar{\omega}(b)$ ($\rho = 0.98$, P value $\ll 0.0001$). However, the estimates are also positively biased, especially for longer branch lengths. Branch lengths are also consistently overestimated, as shown in figure 10B. M0 does not account for monotonic changes in

dN/dS along a branch. It may be that the rapid accumulation of nonsynonymous substitutions that immediately follows the $A^h \rightarrow B^h$ transition at all sites causes both ω and branch lengths to be overestimated. As a topic for further research, it might be useful to devise a codon model that can account for the exponential decay in the rate ratio over time to reduce these biases.

Discussion

The mutation–selection framework of Halpern and Bruno (1998) describes a mechanistic substitution process that is more realistic than that implied by commonly used phenomenological codon models. The simplest codon model M0, and its more sophisticated counterparts, implicitly assume that all amino acids have the same fitness save the one currently occupying the site, as was demonstrated in the section titled “M-Series Models Interpreted as Frequency-Dependent Selection under the MutSel Framework”. MutSel, by contrast, permits amino acids to have different fitnesses. This difference between the two approaches has many implications, several of which we explored via theoretical arguments and computer simulations.

We showed in the section titled “Split MutSel Landscapes” that a site-specific MutSel landscape can in theory be split between two amino acids. Such landscapes can lead to patterns consistent with functional divergence. Suppose a site were to evolve in two segregated populations under the site-specific landscape depicted in figure 3 long enough for each population to be fixed at a different amino acid. Further, suppose that each population were to then undergo a 10-fold increase in N_e , so that the site evolves under the MutSel landscape depicted in figure 2. It could happen that one population becomes fixed at T , and the other at E , depending in part on the starting codon for each population. Subsequent changes at other sites and/or in other genes might then canalize this difference as the populations diverge over macro-evolutionary time scales (Pollock *et al.* 2012). By

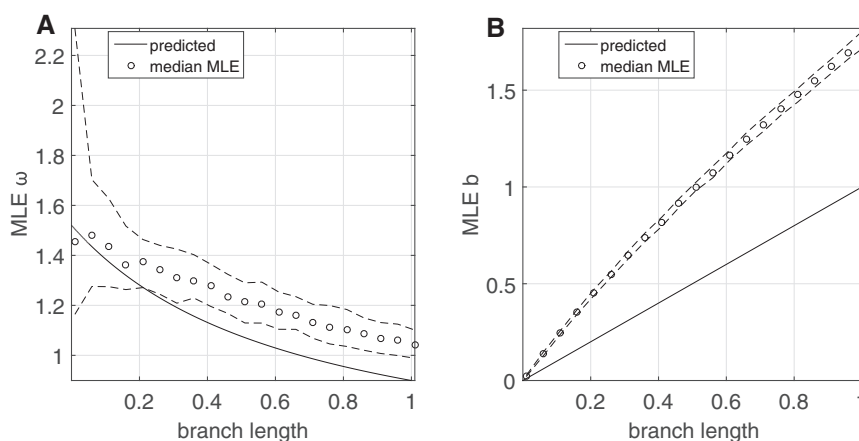


FIG. 10. Circles are median values of maximum likelihood estimates (MLE) produced by fitting data generated under MuSelES to the codon substitution model M0. Dashed lines indicate their inter-quartile range. A: A comparison of the predicted and estimated rate ratios computed from pairs of sequences generated under the MutSelES model. B: A comparison of generating and estimated branch lengths.

this process, a site might eventually exhibit the constant-but-different pattern of Type II divergence without any change in its fitness coefficients.

To cite a real case, Gu (2006) identified sites on the COX gene that exhibited Type II divergence; radical differences in physicochemical properties that are constant within subtrees but different between subtrees. This included a site for which *T* (categorized as hydrophilic) and *E* (charge-negative) dominated the COX1 and COX2 clusters, respectively. Although the change from hydrophilic to charge-negative implies a change in fitness coefficients, our analysis suggests that the observed pattern could have arisen under a static MutSel landscape, and that this is partly due to the fact that *T* and *E* differ by more than one nucleotide (as is frequently the case for pairs of amino acids with different physicochemical properties). Further investigation of split landscapes under the mutation–selection framework might provide insight into additional mechanisms that can lead to the pattern of Type II divergence.

Our analysis of the mechanistic model for shifting balance in the section titled “A Mechanistic Model for Shifting Balance” indicated that sites are relatively free to move across their MutSel landscapes under the scenario where neither selection nor drift dominate. Box plots for p_2^h in figure 5B show that, under the $\sigma = 0.001$ scenario, ~10% of sites are in the tail of their landscapes any one time. If a population evolving under this scenario were to undergo a rapid increase in effective population size, we could expect these sites to be forced toward their peaks in concert to produce a transient $dN/dS > 1$ signature of positive selection similar to the elevation in the rate ratio following an environmental shift (see section titled “MutSel on a Changing Fitness Landscape”), but without changes in fitness. This scenario underlines the potential importance of including changes in effective population size in phenomenological codon substitution models.

Setting aside the effect of changes in N_e , it has been commonly assumed that statistical evidence for $\omega > 1$ at some sites and/or branches in a tree is indicative of positive selection due to episodic or continuous (e.g., frequency-dependent) changes in fitness coefficients. We showed in the section titled “A Mechanistic Model for Shifting Balance” that under certain conditions (e.g., $\sigma = 0.001$, $N_e = 1000$) the substitution process can walk a site over its MutSel landscape in such a way that it frequently moves between its peak and tail, a process we call shifting balance. Shifting balance can manifest as phenomenological switches between $\omega_1 \leq 1$ and $\omega_2 > 1$ that can be detected by commonly used branch-site models, as was demonstrated in the section titled “Detecting Positive Selection Caused by Shifting Balance”. This suggests that we can never be completely certain whether a site inferred to have undergone positive selection did so as a result of changes in fitness coefficients or shifting balance based on statistical analysis of the sequences alone. Additional information about the role of the protein or the history of the organism would no doubt decide the issue in many cases. A protein implicated in an arms race (e.g., an immune surveillance protein in conflict with a pathogenic immune evasion protein) is very likely to have

undergone changes in fitness coefficients at some sites (Hughes and Nei 1988). So too for a protein that has been linked to variations in phenotype correlated with changes in habitat (Yokoyama *et al.* 2008). However, it would seem that, in the absence of corroborating evidence of some kind, shifting balance should be the null hypothesis when ω is inferred to be >1 at some sites and/or branches (especially when branch lengths are no shorter than 0.25 single nucleotide substitutions per codon).

Covariation models were originally devised to capture variations in the rate ratio at a site caused by substitutions at other interacting sites, a phenomenon predicted by the covariation hypothesis (Fitch 1971). Such models maintain the assumption that sites evolve independently to avoid the computational cost associated with modeling epistatic interactions, and aim only to capture site-specific variations in ω . We showed in the section titled “A Mechanistic Model for Shifting Balance” that variations in the rate ratio at a site can occur by another process, shifting balance, and in the section titled “Detecting Transient Changes in ω Caused by Shifting Balance on a Fixed Landscape” that this process is readily detectable by a covariation-like model. These results suggest that it might be necessary to modify covariation-like models to take dependencies between sites into account. A “true” covariation model of this type could potentially be used to distinguish between shifting balance and the covariation process, and thereby provide a means to estimate how prevalent shifting balance is in real alignments.

Although branch-site models commonly used to infer site-specific changes in the rate ratio (CLM3 and BUSTED) can detect shifting balance under some conditions, M-series models that do not allow site-specific variations in ω are insensitive to this process. Spielman and Wilke (2015) proved that the site-specific rate ratio dN^h/dS^h cannot exceed one when synonymous codons are equally fit. Whereas this proof applies to a single site, our empirical results suggest that an equivalent statistical statement holds for the estimate of a single ω for an alignment with variations in dN^h/dS^h across sites. The likelihood ratio test for the contrast of M0 with $\omega = 1$ versus M0 with $\omega > 1$ was applied to all 50 trials in each of our (σ, b) scenarios (supplementary table S2, Supplementary Material online). The test never rejected the null at the 5% level of significance, and so ω was never inferred to be >1 . We showed in the section “M-Series Models Interpreted as Frequency-Dependent Selection under the MutSel Framework” that an M-series model is consistent with frequency-dependent selection among sites for which ω is >1 . We therefore suggest that it may be more appropriate to use M-series models when analyzing a gene suspected to have undergone frequency-dependent selection, if only to remove the possibility of detecting shifting balance (as would be possible using a branch-site model) and confusing it for episodic positive selection.

The analysis of MutSelES in the section titled “MutSel on a Changing Fitness Landscape” underlines the theoretical difference between shifting balance and episodic changes in fitness landscapes. Positive selection by shifting balance is an independent, random, and site-wise process. Some small

proportion of sites will be in the tail of their landscapes and so under positive selection at any given time. But the sites undergoing this process changes over time. By contrast, we assume that an episodic change in environment will cause at least a subset of sites (e.g., sites that correspond to a functional domain or epitope) to undergo changes in their site-specific landscapes at the same time. This difference impacts model performance. Our analysis shows that the signature of positive selection (i.e., $\omega > 1$) under MutSelES is strong enough to be captured by a single ω estimated using M0 provided the branch length is small. By contrast, Spielman and Wilke (2015) showed that M0 will almost always infer $\omega < 1$ when fitness coefficients are fixed (provided synonymous codons are equally fit). However, because CLM3 and BUSTED allow sites to switch rate ratios independently (so that the efficient pruning algorithm can be used to compute likelihoods), they are insensitive to the difference between sites shifting independently and sites switching in concert. This inability to discriminate between the two scenarios might be addressed either by introducing dependence between sites to account for simultaneous changes in landscapes (accepting the computational cost), or by developing *post hoc* analyses that can make the distinction. These would both be interesting subjects for future research.

Shifting balance has implications beyond codon substitution models. For example, consider the method of estimating the proportion α of amino-acid substitutions attributed to positive selection (Smith and Eyre-Walker 2002). Estimates of α range between $\sim 10\%$ for humans to $> 50\%$ for *Drosophila* (Grossmann et al. 2014, and references therein). While substitution by positive selection is often taken as an indication of adaptive evolution, our analysis in the section titled “A Mechanistic Model for Shifting Balance” suggested that as much as 10% of substitutions by positive selection can be attributed to shifting balance on a static fitness landscape. Thus, it seems possible that the human genome might not be evolving in response to changes in selection pressure, but merely experiencing shifting balance. Likewise, some fraction of positive selection in *Drosophila* might be attributable to the same process. The key question therefore is How prevalent is shifting balance in real data? We leave this question open to future efforts.

Methods

Data Generation

All sequence alignments were generated using the mutation–selection framework. Mutations were modeled under HKY for which the mutation rate from nucleotide k to ℓ is:

$$\mu_{k\ell} = \begin{cases} \nu\kappa\pi_{\ell} & \text{if } k \rightarrow \ell \text{ is a transition} \\ \nu\pi_{\ell} & \text{if } k \rightarrow \ell \text{ is a transversion} \end{cases} \quad (18)$$

ν is a scaling parameter set so that $\mu_{k\ell}$ is the expected number of $k \rightarrow \ell$ mutations per generation, and π_{ℓ} is the stationary frequency for nucleotide $\ell \in \{T, C, A, G\}$. We used $\kappa = 4$ and uniform nucleotide frequencies for all of our simulations. Vectors of site-specific fitness coefficients f^h

were drawn from a zero-mean multivariate normal distribution with covariance matrix $\sigma^2 I$, where I is the 61×61 identity matrix. Each vector was modified to make synonymous substitutions equally fit before constructing the site-specific rate matrix. The effective population size was $N_e = 1,000$ for all simulations. Variations in the strength of the shifting balance phenomenon were effected by using values of $\sigma \in \{0.0001, 0.001, 0.01\}$. All alignments were generated on a symmetrical 8-taxon rooted tree with uniform branch lengths $b \in \{0.25, 0.50, 1.00\}$ except where indicated otherwise.

The element A_{ij} of the substitution rate matrix of equation (2) gives the expected number of codon substitutions $i \rightarrow j$ per generation. For purposes of inference however, all A^h were re-scaled so that branch lengths measure the expected number of single nucleotide substitutions per codon. This was done by computing the expected substitution rate $r^h = -\sum_{(i,j)} \pi_i^h A_{ij}^h$ for each rate matrix and then dividing all A^h by the mean $(1/n) \sum_h r^h$. By this re-scaling, the effect of ν in equation (18) is canceled, so ν can remain unspecified. Note that it is possible to normalize the A^h in other ways, for example to make b the expected number of single nucleotide synonymous (neutral) substitutions per codon (Tamuri et al. 2012).

Models M0 and M3

The simplest codon substitution model, referred to as M0 (Yang et al. 2000), assumes that all sites evolve under a single mutation–selection regime characterized by rate ratio ω and transition bias κ . The 61×61 rate matrix Q for this process can be specified as follows for all $i \neq j$:

$$Q_{ij} = \begin{cases} 0 & \text{if } i \text{ and } j \text{ differ by more than one nucleotide} \\ \pi_j & \text{for synonymous transversions} \\ \kappa\pi_j & \text{for synonymous transitions} \\ \omega\pi_j & \text{for nonsynonymous transversions} \\ \omega\kappa\pi_j & \text{for nonsynonymous transitions} \end{cases} \quad (19)$$

where π_i is the stationary frequency for the i^{th} codon. The diagonal elements Q_{ii} are determined by the requirement that the rows of Q sum to zero. This formulation uses codon frequencies. It is also possible to use nucleotide frequencies (Muse and Gaut 1994) by replacing π_j with π_{ℓ} if the transition from codon i to j corresponds to the single nucleotide substitution $k \rightarrow \ell$. Using nucleotide frequencies is consistent the MutSelM0 model described in the section titled “M-Series Models Interpreted as Frequency-Dependent Selection under the MutSel Framework”. M0 provides the basis for mixture models that permit variations in ω across sites, such as the M-series models of Yang et al. (2000). For example, the model designated M3 assumes that sites are distributed across k selection categories $\omega_1, \dots, \omega_k$ at proportions p_1, \dots, p_k that sum to one. The rate matrix for each category is constructed using equation (19) but with its own

ω . We used the variant of M3 with $k = 2$ for the analyses in this article.

The Covarion-Like Model

We define the covarion-like model with two ω categories using a 122×122 rate matrix (cf. Guindon et al. 2004):

$$Q = \frac{1}{r_1} \begin{bmatrix} Q_1 & 0 \\ 0 & Q_2 \end{bmatrix} + \frac{\delta}{r_2} \begin{bmatrix} -p_2 I & p_2 I \\ p_1 I & -p_1 I \end{bmatrix} \quad (20)$$

where Q_k , $k \in \{1, 2\}$, is the 61×61 substitution rate matrix constructed with $\omega = \omega_k$, I is the 61×61 identity matrix and 0 is a matrix of the same size with all entries set to zero. The expected proportion of time a site evolves under ω_k is p_k . The divisor r_1 adjusts the substitution rate so that branch length is the expected number of single nucleotide substitutions per codon. The divisor r_2 adjusts the switching rate so that δ is the expected number of switches per unit branch length. Other authors of covarion-like models have neglected to scale the switching matrix in this way (Guindon et al. 2004); this is the main reason we used our own covarion-like model instead of *fitmodel*. The divisors r_1 and r_2 were computed as follows:

$$r_1 = \sum_{(\text{rows}, \text{cols})} D \begin{bmatrix} Q_1 & 0 \\ 0 & Q_2 \end{bmatrix} \circ \begin{bmatrix} I_N + I_S & 0 \\ 0 & I_N + I_S \end{bmatrix} \quad (21)$$

$$r_2 = \sum_{(\text{rows}, \text{cols})} D \begin{bmatrix} -p_2 I & p_2 I \\ p_1 I & -p_1 I \end{bmatrix} \quad (22)$$

where \circ is the element-wise or Hadamard product. D is the diagonal matrix whose entries are the stationary frequencies for the 122 (codon, ω) state pairs given by the vector $\langle \pi_1 p_1, \dots, \pi_{61} p_1, \pi_1 p_2, \dots, \pi_{61} p_2 \rangle$. Summations are across all rows and columns.

We designate this model CLM3 because it is the same as M3 (Yang et al. 2000) with two ω -categories but with covarion-like switching. The variant CLM3a sets $\omega_1 < 1$ and $\omega_2 = 1$. This is nested in CLM3b for which $\omega_2 > 1$. The M3-CLM3 contrast tests for shifting ($\delta > 0$). Although the two models differ by one parameter (δ), CLM3 is the same as M0 only when $\delta = 0$, on the boundary of the parameter space. The distribution of the log-likelihood ratio for the M3-CLM3 contrast is therefore an equal mixture of a point-mass at zero and a χ_1^2 (e.g., Case 5 in Self and Liang 1987). The CLM3a-CLM3b contrast provides a test for positive selection. The test was conducted only if the maximum likelihood estimate for ω_2 under CLM3 was > 1 and δ was significantly > 0 . The distribution for the log-likelihood ratio for this contrast is χ_1^2 .

The BSREL Model and BUSTED

The branch-site random effects model or BSREL (Kosakovsky Pond et al. 2011) was motivated by the limitations of earlier branch-site models that allowed positive selection to occur at sites only on a prespecified portion of the tree (Yang et al. 2005). In its most general form BSREL allows each site h to have its own distribution of ω across branch b . This is

specified by a vector of branch-and-site-specific rate ratios $\langle \omega_1^{hb}, \dots, \omega_k^{hb} \rangle$ and their concomitant proportions $\langle p_1^{hb}, \dots, p_k^{hb} \rangle$ that sum to one. The transition probability matrix for a site on a branch of length t is computed as the weighted average:

$$p^{hb}(t) = \sum_{i=1}^k p_i^{hb} \exp(Q_{\omega_i^{hb}} t) \quad (23)$$

where $Q_{\omega_i^{hb}}$ is the rate matrix for M0 with $\omega = \omega_i^{hb}$. This approach obviates the need to specify which ω_i^{hb} applies to which branch by allowing the rate ratio for the site to vary randomly across branches. The branch-site unrestricted statistical test for episodic diversification or BUSTED (Murrell et al. 2015) is a simplified version of BSREL. It estimates three rate ratios along with their proportions to provide a common distribution for all sites. The unrestrained model allows each of the three rate ratios to take on any nonnegative value. If the largest is > 1 then a likelihood ratio test for positive selection is conducted. The distribution of the likelihood ratio statistic for this test is an unknown mixture of χ_0^2 , χ_1^2 and χ_2^2 . To be conservative, BUSTED uses χ_2^2 to compute P -values.

Visualizing the MutSel Landscape in Two Dimensions

Suppose A is the MutSel rate matrix for a site, and that a substitution to the i^{th} codon occurs at that site at time $t = 0$. Then the temporal evolution of the 1×61 vector of expected frequencies $\pi(t|i)$ can be characterized as follows (following McCandlish 2011):

$$\pi(t|i) = \pi + \pi(0|i) \sum_{k=2}^{61} \lambda_k^t D^{-1/2} \mathbf{e}_k \mathbf{e}_k^T D^{1/2} \quad (24)$$

D is the 61×61 diagonal matrix whose entries are the site-specific stationary frequencies π ; $\{\mathbf{e}_k\}_{k=2}^{61}$ and $\{\lambda_k\}_{k=2}^{61}$ are the eigenvectors and eigenvalues for the matrix $D^{1/2} P D^{-1/2}$ sorted so that $\lambda_2 \geq \lambda_3 \geq \dots \geq \lambda_{61}$ where $P = \exp(A)$ is the transition probability matrix for the site for a branch of length one; and $\pi(0|i)$ is a vector of zeros but with one at the i^{th} entry (i.e., the initial distribution with a point mass at the i^{th} codon).

Equation (24) shows that the distribution $\pi(t|i)$ of codons at time t converges to the vector of stationary frequencies π for the site, and that the departure from π at any time t can be written as a function of the elements of the eigensystem of $D^{1/2} P D^{-1/2}$. This eigensystem will typically be dominated by \mathbf{e}_2 and \mathbf{e}_3 , the vectors with the two largest eigenvalues (\mathbf{e}_1 corresponds to π). A plot of the components of the transformed eigenvectors:

$$\mathbf{u}_k = (1/\sqrt{1 - \lambda_k}) D^{-1/2} \mathbf{e}_k \text{ for } k \in \{1, 2\} \quad (25)$$

can be used to depict a 2-dimensional approximation of the dynamics captured by equation (24). Each coordinate $\langle \mathbf{u}_2(i), \mathbf{u}_3(i) \rangle$ gives the location of the i^{th} codon in the 2-dimensional landscape. The transformation \mathbf{e}_k to \mathbf{u}_k makes the length of the edges connecting any two points approximately proportional to the expected number of single nucleotide substitutions

required for the site to move from one point to the other and back again along any possible pathway. Note that McCandlish (2011) used the Moran process model (Moran 1958) to construct rate matrices, in which case time is measured by generations rather than single nucleotide substitutions. See McCandlish (2011) for this and other details.

Supplementary Material

Supplementary data are available at *Molecular Biology and Evolution* online.

Acknowledgments

Funding for this research was provided by the Nova Scotia Graduate Scholarship (C.T.J.), and the Natural Sciences and Engineering Research Council of Canada (J.P.B. and E.S.). We would like to thank the reviewers whose input greatly improved the quality of this work.

References

- Bazykin GA. 2015. Changing preferences: deformation of single position amino acid fitness landscapes and evolution of proteins. *Biol. Lett.* 11:1–7.
- Bielawski JP, Dunn KA, Sabehi G, Bèjà O. 2004. Darwinian adaptation of proteorhodopsin to different light intensities in the marine environment. *PNAS* 101:14824–14829.
- dos Reis M. 2013. Available from: <http://arxiv:1311.6682v1> [cited 2013 Nov 26].
- dos Reis M. 2015. How to calculate the non-synonymous to synonymous rate ratio of protein-coding genes under the Fisher-Wright mutation-selection framework. *Biol. Lett.* 11:1–4.
- Fitch W. 1971. The nonidentity of invariable positions in the cytochrome c of different species. *Biochem Genet.* 5:231–241.
- Goldman N, Yang ZH. 1994. Codon-based model of nucleotide substitution for protein-coding dna-sequences. *Mol Biol Evol.* 11:725–736.
- Grossmann TI, Waxman D, Eyre-Walker A. 2014. Fluctuating selection models and McDonald-Kreitman type analyses. *PLoS One* 9:e84540.
- Gu X. 2006. A simple statistical model for estimating type-II (cluster-specific) functional divergence of protein sequences. *Mol Biol Evol.* 23:1937–1945.
- Guindon S, Rodrigo AG, Dyer KA, Huelsenbeck JP. 2004. Modeling the site-specific variation of selection patterns along lineages. *PNAS* 101:12957–12962.
- Halpern AL, Bruno WJ. 1998. Evolutionary distances for protein-coding sequences: modeling site-specific residue frequencies. *Mol Biol Evol.* 15:910–917.
- Hasegawa M, Yano T, Kishino H. 1984. A new molecular clock of mitochondrial dna and the evolution of hominids. *Proc Jpn Acad.* 60:95–98.
- Hughes AL, Nei M. 1988. Pattern of nucleotide substitution at major histocompatibility complex class-I loci reveals overdominant selection. *Nature* 335:167–170.
- Kimura M. 1962. On the probability of fixation of mutant genes in a population. *Genetics* 47:713–719.
- Kosakovsky Pond SL, Murrell B, Fourment M, Frost SDW, Delpont W, Scheffler K. 2011. A random effects branch-site model for detecting episodic diversifying selection. *Mol Biol Evol.* 28:3033–3043.
- Kryazhimskiy A, Plotkin JB. 2008. The population genetics of dn/ds. *PLoS Genet.* 4:e1000304.
- Liberles DA, Teufel AI, Liu L, Stadler T. 2013. On the need for mechanistic models in computational genomics and metagenomics. *Genome Biol Evol.* 5:2008–2018.
- Lu A, Guindon S. 2013. Performance of standard and stochastic branch-site models for detecting positive selection among coding sequences. *Mol Biol Evol.* 31:484–495.
- McCandlish DM. 2011. Visualizing fitness landscapes. *Evolution* 65:1544–1558. doi:10.1111/j.1558-5646.2011.01236.x.
- Moran PAP. 1958. Random processes in genetics. *Math Proc Camb Philos Soc.* 54:12. doi:10.1017/S0305004100033193.
- Mugal CF, Wolf JBW, Kaj I. 2013. Why time matters: codon evolution and the temporal dynamics of dn/ds. *Mol Biol Evol.* 31:212–231.
- Murrell B, Weaver S, Smith MD, Wertheim JO, Murrell S, Aylward A, Eren K, Pollner T, Martin DP, Smith DM, et al. 2015. Gene-wide identification of episodic selection. *Mol Biol Evol.* 32:1365–1371.
- Muse SV, Gaut BS. 1994. A likelihood approach for comparing synonymous and nonsynonymous nucleotide substitution rates, with applications to the chloroplast genome. *Mol Biol Evol.* 11:715–724.
- Mustonen V, Lässig M. 2009. From fitness landscapes to seascapes: non-equilibrium dynamics of selection and adaptation. *Trends Genet.* 25:111–119.
- Nielsen R, Yang Z. 2003. Estimating the distribution of selection coefficients from phylogenetic data with applications to mitochondrial and viral dna. *Mol Biol Evol.* 20:1231–1239.
- Pegueroles C, Laurie S, Albà MM. 2013. Accelerated evolution after gene duplication: a time-dependent process affecting just one copy. *Mol Biol Evol.* 30:1830–1842.
- Pollock DD, Thiltgen G, Goldstein RA. 2012. Amino acid coevolution induces an evolutionary Stokes shift. *PNAS* 109:E1352–E1359.
- Rodrigue N, Philippe H. 2010. Mechanistic revisions of phenomenological modeling strategies in molecular evolution. *Trends Genet.* 26:248–252.
- Ross SM. 1996. Stochastic processes. New York: Wiley.
- Self SG, Liang KY. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio test under nonstandard conditions. *JASA* 82:605–610.
- Sella G, Hirsh AE. 2005. The application of statistical physics to evolutionary biology. *PNAS* 102:9541–9546.
- Smith MD, Wertheim JO, Weaver S, Murrell B, Scheffler K, Pond SLK. 2015. Less is more: an adaptive branch-site random effects model for efficient detection of episodic diversifying selection. *Mol Biol Evol.* 32:1342–1353.
- Smith NGC, Eyre-Walker A. 2002. Adaptive protein evolution in drosophila. *Nature* 415:1022–1024.
- Spielman S, Wilke CO. 2015. The relationship between dn/ds and scaled selection coefficients. *Mol Biol Evol.* 34:1097–1108.
- Swanson WJ, Nielsen R, Yang QF. 2003. Pervasive adaptive evolution in mammalian fertilization proteins. *Mol Biol Evol.* 20:18–20.
- Tamuri AU, dos Reis M, Goldstein RA. 2012. Estimating the distribution of selection coefficients from phylogenetic data using sitewise mutation-selection models. *Genetics* 190:1101–1115.
- Wright S. 1982. The shifting balance theory and macroevolution. *Annu Rev Genet.* 16:1–19.
- Yang ZF, Wang YF, Zhou Y, Gao QS, Zhang EY, Zhu L, Hu YY, Xu CW. 2013. Evolution of land plant genes encoding l-ala-d/l-glu epimerases (aes) via horizontal gene transfer and positive selection. *BCM Plant Biol.* 13:34–43.
- Yang ZH, Bielawski JP. 2000. Statistical methods for detecting molecular adaptation. *Trends Ecol Evol.* 15:496–503.
- Yang ZH, Nielsen R. 1998. Synonymous and nonsynonymous rate variation in nuclear genes of mammals. *J Mol Evol.* 46:409–418.
- Yang ZH, Nielsen R, Goldman N. 2000. Codon-substitution models for heterogeneous selection pressure at amino acid sites. *Genetics* 155:431–449.
- Yang ZH, Wong SWS, Nielsen R. 2005. Bayes empirical bayes inference of amino acid sites under positive selection. *Mol Biol Evol.* 22:1107–1118.
- Yokoyama S, Tada T, Zhang H, Britt L. 2008. Elucidation of phenotypic adaptations: molecular analysis of dim-light vision proteins in vertebrates. *PNAS* 105:13480–13485.