



Article

Juan Guerrero Montero*, Andres Karjus, Kenny Smith and
Richard A. Blythe

Reliable detection and quantification of selective forces in language change

<https://doi.org/10.1515/cllt-2023-0064>

Received June 2, 2023; accepted October 13, 2023; published online October 30, 2023

Abstract: Language change is a cultural evolutionary process in which variants of linguistic variables change in frequency through processes analogous to mutation, selection and genetic drift. In this work, we apply a recently-introduced method to corpus data to quantify the strength of selection in specific instances of historical language change. We first demonstrate, in the context of English irregular verbs, that this method is more reliable and interpretable than similar methods that have previously been applied. We further extend this study to demonstrate that a bias towards phonological simplicity overrides that favouring grammatical simplicity when these are in conflict. Finally, with reference to Spanish spelling reforms, we show that the method can also detect points in time at which selection strengths change, a feature that is generically expected for socially-motivated language change. Together, these results indicate how hypotheses for mechanisms of language change can be tested quantitatively using historical corpus data.

Keywords: language change; evolutionary dynamics; selection; corpus analysis; verb regularisation; spelling reforms

1 Introduction

Human languages undergo constant change as a result of innovations in form and meaning, and competition between new and old forms of expression. For example, a

***Corresponding author: Juan Guerrero Montero**, School of Physics and Astronomy, University of Edinburgh, Edinburgh, Scotland, E-mail: j.a.guerrero-montero@sms.ed.ac.uk. <https://orcid.org/0009-0006-2255-943X>

Andres Karjus, ERA Chair for Cultural Data Analytics, Tallinn University, Tallinn, Estonia; School of Humanities, Tallinn University, Tallinn, Estonia; and Estonian Business School, Tallinn, Estonia

Kenny Smith, Centre for Language Evolution, University of Edinburgh, Edinburgh, Scotland

Richard A. Blythe, School of Physics and Astronomy, University of Edinburgh, Edinburgh, Scotland

phoneme may start being pronounced in a different way, or a new word order may be introduced. A wide range of factors may be responsible for innovation (Walker 2010). These include expressivity, for example, a desire to be noticed, recognisable, amusing, or charming (Keller and Nerlich 1994), and economy, minimizing the effort needed to communicate without compromising the listener's understanding, which may lead to the development of novel, simpler forms (Zipf 1949). Meanwhile, competition mediated by interactions with other components of the language may favour a more consistent mapping between form and function. Alongside these, social factors like prestige or taboo (Labov 2001), may make certain variant forms more or less attractive to certain language users.

In this work, our aim is to quantify the competition between linguistic variants that are available to a speech community, and therewith gain insights into its origins. We achieve this by viewing language change as a cultural evolutionary process (Atkinson and Gray 2005; Croft 2000; Mufwene 2001; Pagel 2009). When modelling cultural evolution (Boyd and Richerson 1988; Cavalli-Sforza and Feldman 1981), it has long been recognised that changes in variant frequencies may arise both from systematic biases (which we refer to generically as *selection*) and random *drift*. While drift may refer to directional change in linguistics following Sapir (1921), we use it here in the cultural evolutionary sense, denoting unbiased stochastic change. Typically one is most interested in identifying the selective forces that cause one variant to be favoured over another, including linguistic (Labov 1994) or social (Labov 2001) factors. Eliminating the possibility that changes may be entirely due to drift is a necessary first step in this endeavour. Initial attempts to achieve this in the context of cultural evolution involved establishing statistical properties of drift and comparing with the corresponding features of empirical data. For example, the distributions of baby names (Hahn and Bentley 2003) and Hittite ceramic bowl types (Steele et al. 2010), as measured at a single point in time, were found to be consistent with the predictions of drift. Under closer examination, however, deviations from drift were found in both cases, for example, by appealing to the rate at which the most abundant types are replaced (Acerbi and Bentley 2014).

Cultural and linguistic datasets provide a potentially rich source of data to constrain parameters in a model of the evolutionary process. In particular, by combining observations of token frequencies at multiple time points, one should achieve greater inferential power than can be achieved by considering only a single point in time. Although such analyses are challenging to construct, a number of forward steps have been made in recent years. For example, the evolution of pottery styles was investigated by appealing to predictions for the number of types remaining after a given time under drift (Kandler and Shennan 2013) and by using simulated trajectories of variant frequencies in an Approximate Bayesian Computation scheme (Kandler and Powell 2015).

Here, we analyse changes in linguistic corpus data with a method based on the Wright–Fisher model of evolution (Fisher 1930; Wright 1931). Although introduced as a model for changes in gene frequencies through biological reproduction, the Wright–Fisher model is also relevant to cultural evolution (Cavalli-Sforza and Feldman 1981). In the specific context of language change, the Wright–Fisher model has been shown to be equivalent to a variety of different conceptual approaches. For example, a mathematical formulation of Croft’s descriptive theory of utterance selection (Croft 2000), itself grounded in Hull’s (2010) generalised analysis of selection, was shown to have the same structure as the Wright–Fisher model (Baxter et al. 2006). Moreover, Reali and Griffiths (2010) showed that a version of the Wright–Fisher model that includes innovation and drift is equivalent to a model of iterated learning where language learners apply Bayesian inference to estimate a variant’s frequency in their linguistic input. Other theories of language change, for example those that invoke a competition between multiple candidate grammars (Yang 2002), can also be viewed as instances of Hull’s generalised analysis of selection, and it has been argued that these may also be represented as a Wright–Fisher model (Blythe and Croft 2021).

The essence of the analysis presented below is to determine the values of parameters in the Wright–Fisher model that maximise the probability that the model generates the series of variant frequencies obtained from a historical corpus. As we set out in Section 2 below, one of these parameters quantifies the strength of selection, and the other the scale of fluctuations arising from random contributions to language change. A difficulty with the Wright–Fisher model is that the mathematical formulæ that describe the evolution are difficult to work with. In genetics, a great deal of effort has been invested in devising reliable approximations that facilitate application to empirical time series (Tataru et al. 2016), an effort that we utilise here in the cultural evolutionary context. Specifically we build on a Beta-with-Spikes approximation (Tataru et al. 2015) in a way that facilitates an efficient and reliable estimation of model parameters, as judged by benchmarking with both real and synthetic data (Montero and Blythe 2023).

In Section 3 we apply this method to historical corpus data in three separate investigations. First, we revisit the set of English verbs with irregular past tense forms that were previously examined by Newberry et al. (2017), Karjus et al. (2020), and Karsdorp et al. (2020), showing that our method is more reliable than that based on a normal approximation of the Wright–Fisher model (Newberry et al. 2017) while offering greater interpretability than a neural-network based time-series classifier (Karsdorp et al. 2020). In common with Newberry et al. (2017), we find that some verbs appear to be irregularising over time.

By itself, the inferred strength of selection is not necessarily informative as to its underlying cause. Our second investigation demonstrates one approach by which

such information can be gleaned. Specifically, we divide English verbs into two sets: those whose regular past tense form contains a repeated consonant, and those that do not. The former set is then subject to a conflict between the greater grammatical simplicity that would be gained by following the regular pattern and the greater phonological simplicity afforded by omitting the repeated consonant (Leben 1973; Stemberger 1981). By comparing the selection strengths between the two sets, we can show that the latter constraint tends to override the former in the context of English verbs.

Finally, we turn to a set of Spanish words that were affected by orthographical reforms in the 18th and 19th centuries. Here, we demonstrate that an unsupervised maximum-likelihood analysis can pinpoint with good accuracy the time at which the reforms were introduced and furthermore quantify the impact of the reform on the linguistic behaviour of the speech community. These last results illustrate that, even with time-series comprising a few measurement points, we can uncover social changes that might not otherwise be apparent. We discuss such opportunities, along with limitations of our method, further in Section 4.

2 Methodology

2.1 Maximum likelihood estimation methods

Maximum likelihood estimation is a conceptually simple yet powerful technique for estimating parameter values in a model and selecting between multiple candidate models. The basic setup is that we have both some empirical data, denoted X , and a probabilistic model that tells us how likely the observation X is given some choice of parameters Θ . We then estimate the values of the parameters by determining the combination that maximises the likelihood of the data. This procedure lies at the heart of many statistical methods, including linear regression. In such a model, parameters are chosen to maximise the likelihood of the data given a statistical model of the residuals (Severini 2000; Silvey 1970), for example, that the residuals are drawn from a normal distribution. It can also be viewed as a special case of Bayesian inference with a uniform prior.

In this work we are concerned with frequency time series, that is, a sequence of measurements $X = \{(x_t, t)\} = \{(x_1, t_1), (x_2, t_2), \dots, (x_m, t_m)\}$ where x_i is the fraction of instances of use of a linguistic variable (e.g., all past-tense forms of a specific verb) in which a particular variant (e.g., the regular form) was used during a short time window centred on time t_i . Thus, the dataset $X = \{(0.2, 1), (0.5, 2), (0.75, 5)\}$ would imply a proportion of usage of the regular form of 20 % at time 1, 50 % at time 2 and 75 % at time 5 (and no frequency data at any other time).

The underlying evolutionary model of language change determines a set of transition probabilities, $\text{Prob}(x_{i+1}, t_{i+1}|x_i, t_i, \Theta)$, that tell us how likely it is that, given a proportion x_i at time t_i and parameters Θ , the proportion will be x_{i+1} at time t_{i+1} . In the previous example, the dataset X would determine the transition probabilities $\text{Prob}(0.5, 2|0.2, 1, \Theta)$ and $\text{Prob}(0.75, 5|0.5, 2, \Theta)$, whose exact numerical values would depend on the choice of model parameters Θ . We assume that contributions to changes in variant frequencies at different points in time are uncorrelated, which means that we can write the likelihood of the entire frequency time series as the product of the transition probabilities for each interval:

$$L(X|\Theta) = \prod_{i=1}^{m-1} \text{Prob}(x_{i+1}, t_{i+1}|x_i, t_i, \Theta). \quad (1)$$

It is this likelihood function that we will maximise to determine the set of parameters Θ that best describes the cultural evolutionary dynamics, and that we will use to compare different models.

There are two main ways to choose the form of the transition probabilities $\text{Prob}(\cdot, \cdot, \Theta)$, a choice that is crucial to parameter estimation and subsequent interpretation. One could simply assume that the frequencies follow a prescribed trajectory between the two points, subject to some fluctuations around them. For example, in linear regression, frequencies would be assumed to vary linearly with time with normally-distributed residuals. Logistic regression is similar, but instead assumes that frequencies vary following a nonlinear logistic function that is commonly used as a model for S-shaped language change (Kauhanen and Walkden 2018; Kroch 1989; Reali and Griffiths 2010). A weakness of this approach is that without an underlying model of language production and transmission that may be operationalised as frequency time series, it is difficult to relate the parameters obtained to the behaviour of individuals or speech communities.

The alternative is to derive the transition probabilities from an explicit agent-based model of language change, many of which can be understood as a variant of the Wright–Fisher model of evolution (Blythe and Croft 2021). As noted in the introduction, the transfer of a model from genetics is justified on theoretical grounds (Croft 2000; Hull 2010) and one can interpret the parameters by appealing to models of language use (Baxter et al. 2006) or iterated Bayesian learning (Reali and Griffiths 2010). A drawback of the Wright–Fisher model is that exact expressions for the transition probabilities (Crow and Kimura 1970) are complex and difficult to work with computationally, as their associated transition matrices may become numerically intractable (Paris et al. 2019). This has motivated many different approximation schemes (Tataru et al. 2016). In this work, we apply a self-contained Beta-with-Spikes approximation scheme that was developed and tested by Montero and Blythe (2023) and found to provide reliable estimates for parameter values without incurring

undue computational cost. In the following we overview the conceptual components of this approach that are most relevant to linguistic applications, directing the reader to Montero and Blythe (2023) for technical details.

2.2 The Wright–Fisher model

The Wright–Fisher model describes a population of N replicating individuals of different types, each of which directly corresponds to a different variant of a linguistic variable. While any number of mutually exclusive types can be included in the model, we will focus on the case with two distinct types. Its extension to three or more distinct types is straightforward (Tataru et al. 2016). The quantity x_t is the proportion of individuals of a specific variant in generation t , as described above. The process of replication has the effect that, in generation $t + 1$, each individual is of the variant of interest with probability $g(x_t, s)$, which depends both on the composition of the population and a measure of selection strength that we denote s . It is assumed that each of the N individuals in the new generation has its type assigned independently. This replication process is illustrated in Figure 1.

The two evolutionary parameters, N and s , can be afforded a linguistic interpretation as follows. The effective population size, N , quantifies the scale of fluctuations in the variant frequencies around a smooth trajectory of change. The smaller the effective population size, N , the larger the fluctuations. When selection is weak (s is close to zero), the time taken for a variant to go extinct is proportional to N (Crow and Kimura 1970). Its interpretation as a population size, effortless in population genetics, does not work as well when studying language change. Through agent-based models of language learning and use (Baxter et al. 2006; Blythe and Croft 2021), we understand that N generically correlates with the size of the speech community. However, heterogeneous social network structures can result in N correlating only weakly with the size of the human population (Blythe and Croft 2021; Bromham et al. 2015; Wichmann et al. 2008). The population size N is also related to the total usage of all variants of a linguistic variable in a given generation, and how long it is retained in memory (Baxter et al. 2006; Real and Griffiths 2010). Intuitively, speakers will be




Figure 1: Schematic representation of the transition from generation t to generation $t + 1$ in a Wright–Fisher process with $N = 10$.

more consistent in their usage of specific variants the more they encounter them, and the longer they recall these encounters. This increased individual consistency will be reflected as smaller fluctuations in time series data. Although our analysis will not allow these different contributions to N to be distinguished, we will be able to determine which linguemes are subject to greater or lesser uncertainty in transmission between speakers.

The selection strength parameter, s , represents a tendency for the variant of interest to increase in relative usage ($s > 0$) or decrease ($s < 0$). Here, s subsumes all factors that could lead to a variant systematically increasing or decreasing in frequency over time, whether they originate in cultural, cognitive or language internal factors (Croft 2000; Labov 1994, 2001, 2010). Similarly to the various factors that may influence the effective population size, we will not be able to distinguish them from the value of s alone. However, as we show below, we can gain useful information by looking for common features of variants which are found to have similar selection strengths.

The parameter s specifies the probability $g(x, s)$ that an individual in generation $t + 1$ is an offspring of an individual with frequency x and selection strength s in the previous generation. Here, we take

$$g(x, s) = \frac{1}{1 + \frac{1-x}{x}e^{-s}}, \quad (2)$$

which has been commonly used in the theoretical characterisation of language change (Kauhanen and Walkden 2018; Yang 2000). In Figure 2 we plot the transition probability $\text{Prob}(x|x_0, s)$ that results from this definition for the case where $x_0 = \frac{1}{2}$ and for different values of s . We see that larger values of s shift the peak of this

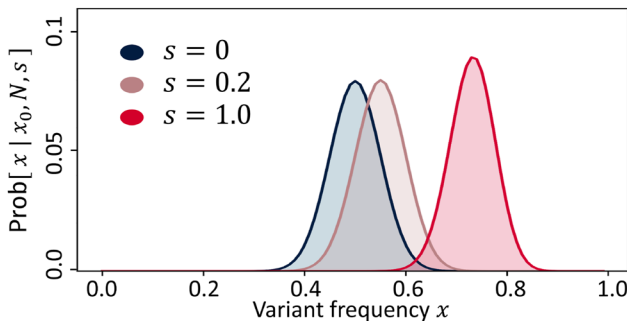


Figure 2: Probability distribution of a variant frequency x after one generation of evolution in the Wright-Fisher model, starting from a frequency $x_0 = \frac{1}{2}$. As the selection strength s increases, the distribution becomes more sharply peaked on larger values of x .

distribution towards higher values of this frequency x , consistent with the notion of a bias towards the corresponding variant.

In the literature, one can find relationships between the selection strength s and the probability $g(x, s)$ different to that specified above (Crow and Kimura 1970; Tataru et al. 2016). Our chosen formula has the useful property that $g(x, s) + g(1 - x, -s) = 1$, which means that if one of two variants in a population has a selection strength s , the other one implicitly has a selection strength $-s$. This choice thus lends a symmetry between positive and negative selection strengths of the same magnitude, which aids the interpretation of the results. The strength $s = 0$ represents pure drift, where any changes in usage over time are due to the stochasticity of replication alone, and not the presence of selective forces. Under pure drift, $g(x, s)$ reduces to $g(x, 0) = x$.

2.3 Beta-with-Spikes approximation

For a single generation of evolution in the Wright–Fisher model, the transition probability is the binomial distribution

$$\text{Prob}(x_{t+1}|x_t, N, s) = \binom{N}{Nx_{t+1}} g(x_t, s)^{Nx_{t+1}} (1 - g(x_t, s))^{N(1-x_{t+1})} \quad (3)$$

because there are N individuals and a success probability of $g(x_t, s)$. If the effective population size N is known, and the time between two frequency measurements corresponds to a single generation, one can use this expression for the transition probability in Equation (1) to construct the overall likelihood of a series of measurements. In the present application to linguistic corpus data, neither of these requirements hold. N is a parameter that we need to estimate, and measurement times are not in general separated by a fixed interval that constitutes a single Wright–Fisher generation. The Beta-with-Spikes approximation, introduced by Tataru et al. (2015) and extended by Montero and Blythe (2023), is designed to deal with these complexities.

For two observations made at times x_t and x_{t+k} (i.e., separated by k generations) the BwS approximation is

$$\begin{aligned} \text{Prob}(x_{t+k}|x_t, N, s) &= P_{0,k} \delta(x_{t+k}) + P_{1,k} \delta(1 - x_{t+k}) \\ &+ (1 - P_{1,k} - P_{0,k}) \frac{x_{t+k}^{\alpha_k - 1} (1 - x_{t+k})^{\beta_k - 1}}{B(\alpha_k, \beta_k)}. \end{aligned} \quad (4)$$

Here, $P_{0,k}$, $P_{1,k}$, α_k and β_k are parameters that determine the shape of the distribution. These parameters have the following interpretation. $P_{0,k}$ is the probability that the variant has gone extinct by the k th generation, and $P_{1,k}$ is the probability that it has

driven the other variant to extinction in that time. α_k and β_k control the shape of the distribution of variant frequencies, conditioned on neither of them having gone extinct. These parameters can be determined from the mean and variance of this conditional distribution (see Montero and Blythe 2023). Note that all four parameters depend on N and s , as well as the sequence of observed frequencies x_{t_i} . Therefore, they need to be recalculated for each time series and combination of model parameters.

A crucial advantage of the BwS approximation is that it accounts for the fact that changes in variant frequencies cannot be arbitrarily large. If a variant has a low frequency (x close to zero), then a downward fluctuation should cause it to become extinct, rather than attain a negative frequency. It is the spikes (represented by the delta functions) in the Beta-with-Spikes expression (4) that incorporate this constraint. By contrast, a normal approximation to the same transition rates (as used by Feder et al. 2014; Newberry et al. 2017) allows, in principle, arbitrarily large or negative x , instead of being constrained to the range $0 \leq x \leq 1$. This difference is illustrated in Figure 3 which shows the statistical distances between the BwS and normal approximation and the exact WF transition probability, for different values of the initial frequency x_0 and two values of the selection strength s . We see that for both pure drift ($s = 0$) and strong selection ($s = 0.5$), the BwS approximation stays consistently closer to the exact distribution for all values of x_0 , which is reflected in lower values of the statistical distance. In particular, the BwS approximation is




Figure 3: Comparison of the statistical distances of the BwS and normal approximations to the exact WF distribution as a function of the initial frequency x_0 . Left: statistical distance for pure drift ($s = 0$). Right: statistical distance for strong selection ($s = 0.5$). The Beta-with-Spikes approximation has lower statistical distance to the exact distribution (meaning it approximates it more accurately) for every value of s and x_0 , but especially for extreme values of x_0 close to 0.0 or 1.0 and for strong selection.

significantly better than the normal approximation for initial frequencies x_0 close to the edges of the interval, and for strong selection.

The main task in applying the BwS approximation is to estimate the parameters $P_{0,k}$, $P_{1,k}$, α_k and β_k for successive generations $k = 1, 2, 3, \dots$. The strategy of Tataru et al. (2015) is to match up the moments of the BwS distribution to those of the Wright–Fisher model after k generations have elapsed. This method works well when the selection strength s is small, but less so when it is large. Montero and Blythe (2023) have improved on the method, particularly in the large s regime, by iterating (3) one generation at a time, and reading off the extinction probabilities, mean and variance at each. The code that implements this procedure, and generates parameter estimates is available here.

In the context of cultural evolution, it is not obvious what period of time counts as a generation in the Wright–Fisher model. In principle, this is a free parameter which would also need to be optimised by maximum likelihood estimation (and furthermore demand an interpretation). Fortunately, this is unnecessary. Montero and Blythe (2023) further show that the optimum values of $1/N$ and s are both proportional to the chosen generation time. In other words, the generation time serves only to set the units in which the parameters N and s are measured. It is however important to use the same generation time across multiple time series when one wishes to compare the values of N and s that are obtained: otherwise, they would be in different units and not comparable. In this work we generally take the shortest time between successive observation points as the generation time. If one makes it shorter than this, the computational effort increases without any improvement in the quality of the estimates obtained. If one makes it longer, one must then aggregate multiple data points which then entails a loss of temporal resolution. However, as we discuss below, it is sometimes beneficial to combine data points to reduce sampling error that is not accounted for in the present maximum likelihood analysis.

2.4 Distinguishing selection from drift

As established in the introduction, the social, linguistic and cognitive forces driving language change are very diverse. Still, their measurable effects can be broadly characterised as belonging to one of two types. Systematic biases drive the evolutionary process in a specific direction, and can be modelled as selective forces. Frequency effects and stochasticity in transmission produce random, unbiased drift whose effects are always present, albeit not always sufficient to explain the behaviour of the data. Quantitative, empirical analyses benefit from the simple yet powerful and flexible characterisation of language change afforded by this binary description.

By using the transition probabilities (4) in the likelihood function (1), we can find the maximum likelihood values of the effective population size N^* and selection strength s^* via

$$(N^*, s^*) = \arg \max L(X|N, s). \quad (5)$$

In practice, we find that the likelihood function $L(X|N, s)$ has a single maximum, which can be located by successively optimising on N at fixed s and vice versa.

It is important to establish whether the selection strength s^* is significantly different to zero: otherwise, the null model of stochastic drift ($s = 0$) would be sufficient to explain the behaviour of the data without the need for selection (Blythe 2012; Newberry et al. 2017). In order to do this, we compare the maximal likelihood under selection, $L(X|N^*, s^*)$, with the maximal likelihood under pure drift. That is, we first restrict to $s = 0$ and determine the optimal effective population size N_0^* :

$$N_0^* = \arg \max L(X|N, 0). \quad (6)$$

Then we compare the models with and without selection by computing the likelihood-ratio

$$\lambda = 2 \ln \left(\frac{L(X|N^*, s^*)}{L(X|N_0^*, 0)} \right). \quad (7)$$

This quantity can be compared to a reference distribution to find a p -value, an estimation of the probability that the observed time series could have arisen from drift alone¹ (Severini 2000; Silvey 1970). To achieve this, we follow the procedure outlined by Feder et al. (2014) and generate 1,000 artificial time series spanning the same time period as the empirical data X with parameter values $s = 0$ and $N = N_0^*$. For each of these we compute the maximum likelihood values N^* , s^* and N_0^* , using the same sequence of steps as for the original empirical time series. We then compute the likelihood ratio λ and determine what fraction of the artificial time series has a larger λ than the one that was observed. This provides an empirical p -value for the null hypothesis of drift.

3 Results

We now apply the methods set out above in three separate tasks, each with a distinct purpose. First, we revisit the set of verb time series from the Corpus of Historical

¹ The commonly used χ^2 distribution does not work well for p -value estimation when working with data from historical corpora, as it only converges to the true distribution of likelihood-ratios for time series of infinite length (Feder et al. 2014).

American English (COHA, Davies 2010) to benchmark our approach against those of Newberry et al. (2017) and Karsdorp et al. (2020). These results demonstrate that the BwS method is both more robust than a similar likelihood-based approach (Newberry et al. 2017) and more informative than a neural network trained to perform a binary classification (Karsdorp et al. 2020). We also introduce a method for assessing the variability of parameter values under different binning strategies, thereby facilitating a judgement as to which results are more robust.

We then perform similar analyses to understand the direction of selection in the context of English irregular verbs, this time using the English 2019 1- and 2-g datasets from the Google Books corpus (Michel et al. 2011). This larger corpus contains more instances of verbs that appear to be irregularising over time. We find that a phonological constraint that disfavors repeated consonants can override a general preference for regularity. Finally, we use data from the 2019 Spanish 1-g Google Books corpus to show that the dates at which Spanish spelling reforms were introduced can be detected using the unsupervised maximum-likelihood analysis.

The validity of using frequency data from the Google Books corpus to draw conclusions on cultural evolution and language change has been questioned by Pechenick et al. (2015) due to the over-representation of scientific literature in the English sub-corpus throughout the 20th and 21st centuries. While they propose restricting studies of cultural and language change to the fiction sub-corpus, we believe that using frequency data from the general English sub-corpus is justified for the purposes of our study. First, our work rests on the comparison between two data sets of English verbs differing only in their phonology. It is reasonable to assume that, if any bias exists in scientific texts regarding the use of irregular or regular forms of verbs, this bias will not be phonologically conditioned, thus maintaining the validity of the comparison between both data sets. Secondly, we have chosen verbs that are reasonably present in both the general English corpus and the English Fiction corpus, so potential biases towards uncommon verbs in scientific literature should not be an issue. Thirdly, the general English sub-corpus will contain more words than the restricted fiction sub-corpus, thus reducing the effect of sampling noise on our results.

3.1 Drift versus selection in past-tense English verbs

A simple example of competition between two variants is provided by English verbs with an irregular past tense form which in many cases coexists with a regular form. This competition has been studied from a variety of quantitative perspectives (Cuskley et al. 2014; Karjus et al. 2020; Karsdorp et al. 2020; Lieberman et al. 2007; Newberry et al. 2017; Ringe and Yang 2022). Of greatest relevance to the present work

are those studies that aimed to distinguish drift from selection as the mechanism behind changes in the relative frequencies of the regular and irregular forms over time.

Newberry et al. (2017) applied the Frequency Increment Test (FIT, Feder et al. 2014) to a set of verbs from the Corpus of Historical American English (COHA, Davies 2010). This is a maximum-likelihood method that rests on a normal approximation to the Wright–Fisher transition probabilities. Like the Beta-with-Spikes maximum-likelihood method in Section 2.4, this method yields estimates of the effective population size and selection strength, along with a p -value for the null hypothesis of pure drift. However, there are situations where results are flagged as unreliable due to the frequency increments failing a normality test (Newberry et al. 2017). Karjus et al. (2020) further noted that the results can also be sensitive to the size of the window over which frequencies are estimated.

Karsdorp et al. (2020) avoid these issues by taking the rather different approach of training a neural network on simulated time series generated by the Wright–Fisher transition probabilities (3) for different values of s (but fixed $N = 1,000$). Each time series in the training set is labelled according to whether it was generated purely by drift ($s = 0$) or if selection was operating ($s \neq 0$). Once trained, the network yields a binary classification of empirical time series, according to whether they are more similar to the examples of drift or selection in the training data. We refer to this as the Time Series Classification (TSC) approach. The advantage of TSC is that no approximation to the Wright–Fisher transition probabilities is made. Moreover, one can manipulate the training data so that it displays artifacts of binning or finite sample sizes that are features of real time series, which in turn should improve the reliability of the classification. This approach does however come with some drawbacks. Whilst the output from the classification algorithm is a value between 0 and 1, it does not have an obvious interpretation as a probability. Karsdorp et al. (2020) used a threshold of 0.5 to label timeseries as arising from drift or selection. The method further does not provide an estimate of the strength of s , and since N was fixed in the training set, this amounts to an assumption that this single value of N was appropriate for all empirical time series. This could be an issue since Newberry et al. (2017) report a wide range of values of N for this data set (from around 80 to around 22,500).

In Section A of Appendix we report the maximum likelihood estimates of N and s , along with the p -value for the drift hypothesis, obtained using the BwS method for the same set of verbs that were considered by Newberry et al. (2017) and Karjus et al. (2020) using FIT and by Karsdorp et al. (2020) using TSC. We perform the analysis by extracting annual frequency data of the variants of interest from COHA and aggregating it into 10-, 20- and 40-year bins. The reason for this is a trade-off between the more precise frequency estimates that derive from larger bins and the greater

temporal resolution obtained from a larger number of bins over the relevant historical period. By employing different binning strategies, we can gain insights into the consequences of this trade-off. Variable-width binning strategies have also been successfully applied in previous studies (Newberry et al. 2017). In these, the number of tokens per bin is kept roughly constant at an arbitrarily chosen value, at the expense of varying their temporal width. For the purpose of comparing the different methods, we have chosen to look only at fixed-binning strategies, although the BwS method could be combined with variable-width binning.

We focus first on the role played by selective forces, which we quantify by appealing to the p values associated with the null hypothesis of pure drift as described in Section 2.4. In Figure 4 we compare the results obtained from the three different methods by ordering the verbs from left to right by decreasing BwS p -value, averaged over the three temporal binnings. Each panel corresponds to a different analysis method, and indicates the p -value for the hypothesis of pure drift for each verb and binning protocol. We recall that higher p -values are more suggestive of the historical changes being due to drift: these are represented with colours ranging from light to dark blue, with darker colours representing higher p -values. Meanwhile, low p -values point towards other forces (such as selection) being present and are represented with different shades of red. While we use the standard p -value threshold

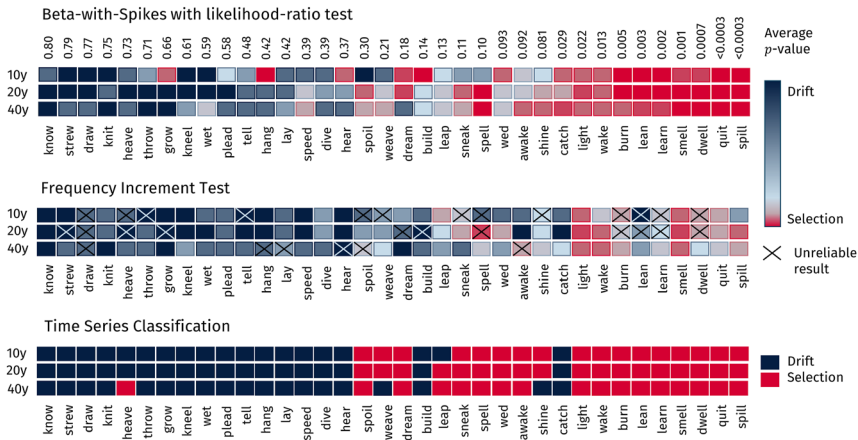


Figure 4: Results for the detection of selective forces in 36 COHA verbs, with three different methods and for three different temporal binnings of 10, 20 and 40 years. Results for both the FIT and BwS likelihood-ratio algorithms produce a p -value for the pure drift hypothesis. Blue shades represent higher p -values (i.e., higher likelihood of the data under drift), while red shades represent p -values under the traditional 0.05 threshold of significance for selection. Time series where the normal approximation that FIT relies on is inaccurate are crossed out. Results for the TSC method from Karsdorp et al. (2020) are classified in a binary way as either drift or selection. The average p -value across the three bins widths obtained through the BwS algorithm is shown along the horizontal axis. We note that the BwS method gives results consistent with TSC when FIT is unreliable.

of 0.05 in the transition between blue (drift) and red (selection) in this representation, we acknowledge that these mechanisms lie in a continuum by making the transition between these extremes smooth.

We see from Figure 4 that the three distinct methods give broadly consistent results, with those verbs towards the left being more compatible with change through pure drift, and those to the right with change from selection. More precisely, the correlation coefficients between the p -values obtained with different methods are 0.63 (Pearson) between FIT and BwS, 0.68 (biserial) between TSC and FIT, and 0.62 (biserial) between BwS and TSC. Analyses producing high p -values for selection (i.e. implying that drift alone can explain the behaviour of the data) are indicated with blue colours, whereas those where selection is more significant are red. Results obtained through the FIT method are generally consistent with those obtained with the BwS method. However, 30 of the FIT results (27.8 % of the total) are flagged as ‘unreliable’ due to a failed normality test. These reliability issues are designed out of the BwS method, as it does not require normally-distributed increments (Montero and Blythe 2023; Tataru et al. 2016). Confidence in the method’s reliability is also gained by benchmarking with synthetic and genetic data (Montero and Blythe 2023) and through the consistency with the independent TSC results. The higher precision of the BwS at high selection strengths leads to higher significance (lower p -values) in its detection of selective forces when compared to the normal approximation, leading to redder colours in Figure 3.

The TSC appears to give a cleaner classification of verbs according to drift and selection, and greater consistency with different choices of bin size. This is likely due in part to the training data being subjected to the same binning protocol as the empirical time-series, but also because a strict threshold was applied to the neural network’s output value to partition into the two classes. While the TSC neural networks produce a value between 0 and 1 as their output, making it more nuanced than this binary classification would suggest, this number is not a probability or a p -value like those produced by BwS or FIT. Thus, an arbitrary threshold is necessary in order to classify time series as driven by drift or selection. A higher or lower threshold would put the boundary between the two classes in a different place. This hinders the interpretability of the result and the estimation of significance levels.

Our results further demonstrate that variation in p -values under different binning strategies, previously observed within the FIT analysis (Karjus et al. 2020), remains evident under the less restrictive BwS analysis presented here. We consequently regard this variability as an inherent feature of the time series data: that is, some changes are harder to classify than others. That is, this uncertainty need not be a failure of the method, but a reflection of linguistic reality. For example, it could reflect different variants being used less predictably by speakers, or by the constraints on variation changing over short timescales (Tagliamonte 2011).

Such observations motivate a more detailed investigation of the classifiability of individual time series. A time series that shows limited variation in parameter values under different temporal binnings is more classifiable than one that shows more variation. With our interest in selection, the two most relevant parameters are s , the selection strength, and the p -value associated with the drift hypothesis. We can visualise the variation in these parameters by performing a Principal Components Analysis (Jolliffe 2002) on combinations obtained through different binning strategies (in this case, bins of 10, 20 and 40 years). The interior of the resulting ellipses indicates the range of variation of the two parameters over different binning strategies. This way, they provide a visualization of not just the average, but the uncertainty and covariance of s and the p -value under different binning strategies. We show these ellipses for the COHA verbs in Figure 5. The upper panel contains the full range of p and s values obtained through the analysis, while the lower panel zooms in on the region where the drift p -value is smaller than 0.05 (i.e., the conventional threshold for rejecting the null hypothesis). We see a correlation between the maximum likelihood value of s and the p -value (both through the positions and rotation angles of each ellipse).

The ellipses that lie entirely within the lower panel correspond to the verbs that are most likely to be driven by selection. We see a clear split between four verbs with positive selection (catch, light, wake and quit), which corresponds to them becoming more irregular over time, and six verbs (learn, lean, burn, smell, dwell and spill) with negative selection, and thus regularising over time. In this analysis, the frequency x is the fraction of irregular forms used in the relevant context. Across the entire plane, there is evidence of both regularisation and irregularisation, although in most cases it is difficult to rule out drift as an explanation for the changes, as was observed by Newberry et al. (2017).

In interpreting these results, it is important to recognise that the presence of fluctuations around a smooth change trajectory will tend towards a higher drift p -value, since in the analysis drift is the sole source of fluctuations. It is possible that fluctuations in the corpus derive from other sources, such as sampling effects associated with a finite corpus. Some methods for estimating parameters in the Wright–Fisher model attempt to account for such fluctuations separately to drift (Tataru et al. 2016). These are, however, typically difficult to implement, and instead we sidestep the issue by ensuring sufficiently many tokens in each temporal bin that the frequency is well estimated. As such, we might expect to see stronger evidence for selection as the bin width is increased, which appears to be true for some (but not all) of the verbs with intermediate p -values. This suggests that some language changes may be dominated by the random effects of drift and therefore exhibit strong fluctuations even in very large corpora.




Figure 5: Variability in the selection strengths s and p -values for the null hypothesis pure drift for the COHA verbs. Each cross shows the mean value of the two parameters for each verb obtained when aggregating frequencies into temporal bins of different lengths. Each ellipse indicates the variability in the parameters at the level of one standard deviation. The vertical axis is an indicator of selection, defined as one minus the p -value associated with the drift hypothesis. The lower panel shows those verbs that fall within the range of p -values that is conventionally used to reject the null hypothesis for a single observation. In this panel we see a clear split into those that are regularising (negative s) and are irregularising (positive s).

To summarise, we have shown in this section that the BwS method can be readily applied to historical corpus data for changes in the frequencies of linguistic variants. It provides estimates of parameters in the Wright–Fisher model that do not rest on an assumption that frequency increments are drawn from a normal distribution, and we find broad consistency in the strength of support for a drift hypothesis with complementary methods.

3.2 Competing linguistic motivations in English verbs

In the previous section we observed a split between some verbs that were regularising and some that were irregularising. While the extension of regular inflection at the expense of irregulars seems to be the norm (e.g. Bybee 1995; Sims-Williams 2016), irregularisation is however an attested phenomenon. Cuskley et al. (2014) found that the processes of regularisation and irregularisation tend to take place with similar frequency, something that is also perhaps suggested by Figure 5, which shows a similar density of verbs along the branch with positive *s* (towards irregularity) and negative *s* (towards regularity). Ringe and Yang (2022) suggest that irregularisation may occur if the number of verbs within an irregularity class is high enough to surpass a productivity threshold. Following Bybee (2001) and Prasada and Pinker (1993), both Cuskley et al. (2014) and Newberry et al. (2017) propose phonological analogy as a potential mechanism for irregularisation. Couched in the terms of the present work, this would correspond to the general rule (adding -ed) contributing a negative value to *s* whilst rules that apply only to a specific subset of verbs contribute a positive value to *s*. Note that we do not necessarily imply that these contributions are additive: for example, in optimality theory (Hayes 1999; Prince and Smolensky 1997), higher-ranked rules take precedence over lower-ranked rules. In general, we may regard opposing forces on linguistic variation as arising from competing motivations which have been discussed in a variety of language change contexts (e.g., Bates and MacWhinney 1987, 1989; DuBois 1985; Haiman 1983; Hawkins 2004; Kirby 1997). By whatever mechanism this opposition is resolved, an overall positive *s* value here indicates that the irregularising rule is dominant.

In this section, we investigate a distinct motivation that may favour irregularisation, namely the phonological simplicity that is afforded by omitting a sound repetition that would occur under application of the regular rule. Specifically, we consider verbs whose infinitives end in alveolar stops (/d/ or /t/) and have an irregular past form where the regular -ed termination is omitted. Examples include *I bled* instead of *I bleded* or *she bet* instead of *she betted*. Verbs where devoicing of final /d/ or changes in the root vowel take place on top of the omission of the termination are also considered. Thus, we hypothesise that the regular form is preferred from the point of view of inflectional simplicity (i.e. using the regular everywhere leads to a simpler inflectional system), while the irregular form is favoured by phonological simplicity. By applying the BwS algorithm to estimate the *s* parameter (and in particular, its sign), we can assess how these competing motivations play out.

For this investigation we switch to the 2019 English Google Books corpus (Michel et al. 2011), as the number of verbs falling into this category and whose past tense forms are both sufficiently frequent and can be reliably identified is relatively small.

The larger size of Google Books relative to COHA allows more examples to be included. We identified 19 English verbs whose irregular and regular forms both show usage above 1% at least in one 5-year bin in the Google Books corpus in the considered time frame (1809–2009). These verbs are: *bend, bet, bite, blend, build, fit, glide, knit, light, pat, plead, quit, slide, speed, spit, thrust, tread, wed, and wet*. A difficulty in the analysis is that the irregular past-tense form can coincide with certain present-tense forms. A major exception is when the verb is preceded by a third-person singular pronoun (e.g., the present *he bets* vs. the irregular past *she bet*), which can easily be distinguished in the bigram dataset. We recognise that this separation is not perfect: for example, certain English varieties do not use the third person marker -s, but we consider the effect of these contributions to be negligible in the corpus. We also kept only those cases where the pronoun was judged to appear at the start of a sentence (by virtue of capitalisation), so as to exclude contexts where the pronoun is followed by the infinitive in a question or an inversion. Again there are situations where capitalised pronouns can appear mid-sentence, but these are also rare. With this, total counts of usage for verbs with non distinct irregular past tense forms range roughly between 2,600 (*knit*) and 120,000 (*pat*), while counts for verbs whose irregular past tense is distinct from their base form range between 600,000 (*glide*) and 40,000,000 (*build*).

In order to formally test whether a potential bias towards irregularisation is significant, a similar analysis was carried out on a baseline set of 34 English verbs whose base form does not end in /d/ or /t/. Data was extracted from Google Books and all verbs satisfy the same conditions on minimal usage in the time frame of interest (1809–2009). The chosen verbs are: *awake, blow, burn, catch, cleave, creep, dive, dream, dwell, freeze, grow, hang, heave, hew, kneel, lean, leap, learn, shake, shear, shine, slay, slink, smell, sneak, spell, spill, spoil, strew, string, strive, swell, wake and weave*. Total usage for these verbs in the Google Books corpus for the specified period ranges between 211,000 tokens (*slink*) and 31,900,000 (*learn*), in the same orders of magnitude as the /d/,/t/ set.

The maximum likelihood parameters for these 53 verbs are given in Section B of Appendix. Here, we visualise our findings by plotting ellipses in the plane spanned by the selection strength and the indicator of selection, following the same procedure as previously described for the COHA verbs, albeit with the addition of a 5-year temporal binning strategy. With this, each ellipse in the *s-p* plane for each verb is produced by averaging the results of the analyses of at most four temporal binnings. We recall that these ellipses characterise the variability in these parameters as the temporal binning is varied. The upper panel in Figure 6 shows the results for all 53 verbs.

For the purpose of comparing the two sets of verbs, we partition the *s-p* plane into four regions: those with positive or negative selection strengths; and those

Table 1: Contingency table for the comparison of irregularising behaviour between the set of verbs ending in alveolar stops and the baseline set. Irregularisation is significantly more common amongst verbs ending in alveolar stops, with a p -value of 0.031 as provided by the G-test.

	Irregularising	Inconclusive	Non-irregularising
Alveolar stop set	9	2	8
Baseline set	7	4	23

plead, swell, thrust) have confidence regions only partially contained in this region of the s - p plane, indicating that, while selective forces towards the irregular form are a plausible explanation to their dynamics, the pure drift hypothesis cannot be confidently ruled out. The remaining 31 verbs (8 in the alveolar stop set, 23 in the baseline set) have confidence regions contained entirely outside this region of likely irregularisation.

These results suggest that verbs in the alveolar stop set are more likely to be selected towards their phonologically simpler irregular form than their counterparts in the baseline set. To test the significance of these findings, we construct the 2×3 contingency table shown in Table 1, where one dimension expresses belonging to the alveolar stop or the baseline sets, while the other dimension expresses whether the verbs' ellipse falls in the irregularisation region in the bottom panels of Figure 6. The p -value for the null hypothesis that the baseline and alveolar stop verbs are drawn from the same distribution is 0.031, as obtained by applying the G-test of goodness-of-fit to the contingency table (McDonald 2014). This indicates that the specific rule favouring phonological simplicity likely outcompetes a general tendency towards regularity.

It is possible that other effects may be responsible for this subset of verbs tending to irregularise. For example, it is well understood that higher frequency items tend to tolerate greater irregularity (Bybee 2007). Given the selection criteria imposed to arrive at the set of 12 verbs in this analysis, it is possible that the sample is skewed towards higher frequency and more irregular forms. However, as noted, the total token counts for both the baseline set and the alveolar stop set span similar ranges, and also have similar averages (of around 5 million for both sets). Therefore we consider this alternative explanation unlikely.

This is not the only phonological conditioning on irregularisation that can be inferred from Figure 6. The subset of verbs ending in a short vowel plus a lateral (*dwel, smell, spell, spill, swell*) seem to be a lot more likely to regularise under selection than other verbs in the study. A similar G-test to the one performed on the alveolar stop test on Table 1 reveals that this tendency is significant with $p < 0.003$. The origin of this tendency is, however, unclear.

To summarise, in this section we have shown that by focussing on a subset of verbs that are subject to specific combination of competing motivations, the Wright–Fisher model combined with the BwS approximation can be used to determine the net effect of this competition. Specifically, we have acquired evidence that phonological simplicity dominates inflectional simplicity in this competition, suggesting perhaps that this is an instance of an OCP constraint (Obligatory Contour Principle, Leben 1973; Stemberger 1981). OCP constraints disfavour pairs of identical or near-identical consonants from being in close proximity to each other. In particular, the constraint here appears to be an OCP-place constraint (Frisch et al. 2004; McCarthy 1986; Pozdniakov and Segerer 2007), meaning that it does not just affect identical consonants, but all alveolar stops independently of voicing.

3.3 Spanish spelling reforms

So far we have assumed that the evolutionary parameters (the effective population size N and the selection strength s) have been constant over time. In the case of competition between regular and irregular verbs this is a reasonable assumption, due to the factors favouring one over the other likely being cognitive or linguistic in origin. By contrast, social pressures like prestige, taboo, or language contact (Hernández-Campoy and Conde-Silvestre 2012; Labov 2001; McMahon 1994) are inherently time-dependent, and we may expect the selection strength in particular to change over time. Here, we investigate this possibility in the context of a purposeful change made by a regulating institution through prescriptive grammar and spelling rules (Anderwald 2012; Rubin et al. 2013), the acceptance or rejection of which we expect to be reflected by a change in the value of s . While well established algorithms like change-point analysis (Taylor 2000) exist for the detection of change in time series, these suffer from shortcomings that make them inadequate for a more nuanced analysis of change in language and culture. First, change-point analysis is based on the assumption that the data is distributed around a constant average before and after a change, which changes the value of said average instantaneously. This makes this methodology only fit for the detection of rapidly occurring S-shaped curves of language change, where the usage frequency of a variant quickly changes and stabilises. Secondly, change-point analysis provides no extra linguistic information, as it does not assume a model of the underlying evolutionary dynamics. Montero and Blythe (2023) solve this issue by setting out a procedure for estimating times at which the parameters N and s change, thus measuring changes in the evolutionary dynamics of the data rather than its average. We briefly recapitulate and then apply this method below.

The specific changes of interest are spelling reforms in Spanish that were introduced by the Real Academia Española (RAE), the central regulatory institution of the standard Spanish language. Since its creation in 1713, the RAE has regulated Spanish orthography following the phonemic principle over etymological or conservative approaches (Baddeley and Voeste 2012). We study words affected by one of the following reforms: (A) The simplification of the <ss> digraph to a single <s> in 1763, due to the different sounds that both spellings represented having merged in the 16th century (Real Academia Española 1763); (B) The replacement in 1815 of etymological <x> with <j> in all non word-final contexts where it represented the phoneme /x/ (Real Academia Española 1815); (C) The replacement, also in 1815, of <y> with <i> in all non word-final closing diphthongs; (D) The reversal of accentuation rules for words ending in <n>, introduced in 1881. This reform stipulated that words ending in <n> with a tonic last syllable had to be accentuated, while words ending in <n> with a tonic penultimate syllable lost their previously prescribed accent (Real Academia Española 1881). We treat words that gain an accent and words that lose an accent as independent sets (D.1 and D.2, respectively).

We now seek to estimate the time at which each reform occurred by appealing only to the time series data and no external information. The basic idea (see also Montero and Blythe 2023) is to allow different parameter combinations (N, s) to apply before and after a time T . That is, for $t < T$ the Wright–Fisher model with parameters (N_1, s_1) applies, and for $t > T$ the parameters (N_2, s_2) apply. The data likelihood, obtained by combining Eqs. (1) and (4), is then maximised with respect to all five parameters (i.e., N and s each before and after the change, and the time T of the change itself).

After identifying the time T that maximises the data likelihood, one needs to determine if the additional complexity of the five-parameter model is compensated by a sufficiently improved description of the data. To achieve this, we obtain an empirical p -value for the null hypothesis that the selection strength s was constant over the entire time period by following a procedure similar to that described in Section 2.4. Specifically, we determine the maximum likelihood values of N and s without a change point, and generate 500 synthetic time series that match the length of the observed series with these parameter values. For each of these time series, we then optimise the five-parameter likelihood that applies when the selection strength changes at a single point in time. An empirical p -value is then given by the fraction of such time series whose five-parameter likelihood exceeds that of the real trajectory. Although computational constraints limit the number of synthetic time series that can be analysed this way, we find that situations where the five-parameter fit has a high likelihood are extremely rare, and there is little to be gained by estimating their rarity to greater precision. One can then apply a threshold, e.g., $p < 0.05$, to decide whether to accept the more complex model. Having split the time series once, one can

apply the method again to each sub-series, thereby identifying secondary change points. This procedure terminates when none of the sub-series admits a subdivision that yields a sufficiently improved description of the data according to the threshold that has been imposed.

To apply this method to the Spanish Spelling reforms, we identify a set of commonly used words that are affected by each one, and average the relative frequencies of usage of their old spellings over all members of each set. The number of words in each set ranges from 16 to 27. The exact sets are specified in Section C of Appendix. This procedure generates a single effective time series for each of the reforms, and has been found effective in related corpus analyses (Amato et al. 2018).

While this averaging over sets of words decreases the sampling noise in the data and increases the inferential power of the analysis, cultural data still suffers from issues that may affect the applicability of the method. Particularly, corpora tend to contain lower token counts in earlier time periods. When translated to frequency time series, this leads to greater sampling noise fluctuations that may be misidentified as changes in the effective population size parameter N . This issue can be remedied by applying a sampling error equalisation algorithm, as laid out by Montero and Blythe (2023). This method creates subsamples of the larger token counts in the data set, in such a way that sampling effects are of equal magnitude throughout the data. In this way, any significant changes in N detected by the method must be due to changes in effective population size parameter, and not a consequence of unequal sampling noise.

Our results are shown in Figure 7. Despite the aggregation of words within each category (to improve the inferential power) and 5-year bins (to reduce computational effort), we find that the resulting trajectories are still subject to considerable fluctuations. The frequency plotted is that of the deprecated variant, which we find is eliminated in all five cases—this highlights the acceptance and influence of the Real Academia Española amongst the literate population. We show with a red line and dot the time at which the reform was introduced, and with a black dot and solid line the first time T at which subdividing the time-series improves the fit to the data, with a p -value threshold of $p = 0.05$ applied.

In all five cases, we find evidence that the selection s changed significantly over time. In each case, the first detected change point falls within twelve years of the reform being introduced, even when the trajectory is strongly fluctuating. We note that the algorithm does have a tendency to detect the reform after it has occurred, rather than at its inception. This is due to the algorithm not distinguishing past from present, making both the beginning and the end of the sharp decline following a reform are considered equivalent.

Further iterating the algorithm, we can further subdivide the time series, as described above. In doing so, we detect secondary time divisions (dashed lines in

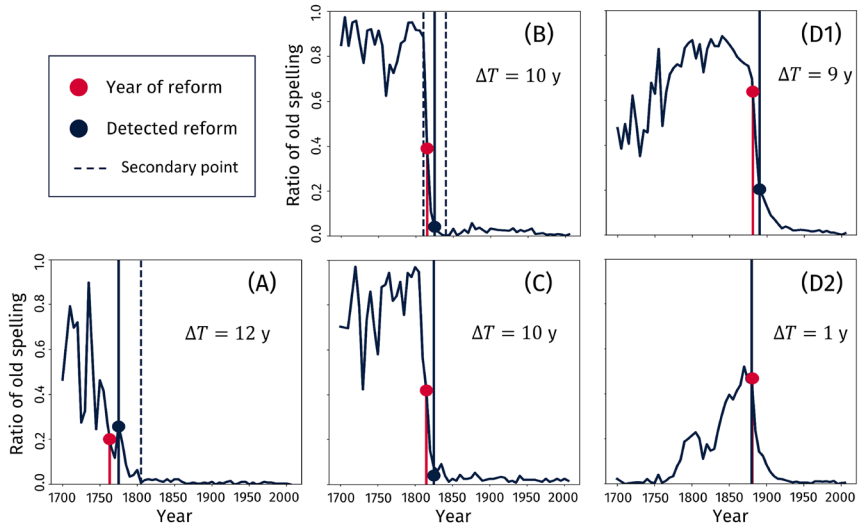


Figure 7: Application of the BwS algorithm for the detection of changing forces to the reference data set of Spanish spelling reforms in the 2019 Spanish 1-g Google Books corpus, with temporal binning of the frequency data of 5 years. For each set of words that undergo a rule change, the ratio of usage of the old form is plotted over time. The ratio of usage of all old forms converges to zero after each reform. Red dots with solid vertical lines represent the year of publication of the RAE spelling reforms (Real Academia Española 1763, 1815, 1881). Dark blue dots with solid vertical lines represent the year at which selection strengths changed as detected by the maximum likelihood method with a p -value below 0.05. These fall within a period ΔT of 12 years or less relative to the date of the reform. Note that the temporal resolution of the time series is of 5 years, so an error of 10 years is equivalent to just two data points. Dashed vertical lines represent secondary points of change in evolutionary parameters, also detected with a p -value below 0.05. The number of such secondary points depends on the time series.

Figure 7), whose p -values are below 0.05. In time series (B), the earlier secondary point detects the beginning of the rapid decline in usage that was deemed less significant than the end by the first application of the algorithm. The later secondary point in (B) and the secondary point in (A) are not associated to documented reforms, and may reflect slight changes in social attitudes or simply be quirks of the data.

Table 2 further records the s -values before and after the main change point. All s , N and p -values for every main and secondary point detected by the algorithm can be found in Section D of Appendix. For categories (A), (D.1) and (D.2), we find that the s value decreases after the detected year of the reform, corresponding to an acceptance of the reform by the speech community. The other two categories however show the opposite trend, with the s value becoming less negative across the reform. We note from Figure 7 that both categories (B) and (C) feature a rapid elimination of the deprecated form, and that this change was in progress before the reform was

Table 2: Maximum likelihood estimates of the first detected time at which the selection strength changed, and its values before and after the change, for the five Spanish reform categories. All changes significant with $p < 0.002$.

	Reform	Detected year	s before division	s after division
(A)	<ss>to <s>	1775	−0.008	−0.37
(B)	<x>to <j>	1825	−0.29	−0.07
(C)	<y>to <i>	1815	−0.30	−0.06
(D1)	Accentuation	1890	−0.05	−0.21
(D2)	Loss of accent	1880	0.13	−0.53

introduced. It has been suggested that in many cases, language reforms tend to reflect pre-existing trends, as opposed to actuating the change (Rutten and Vosters 2021). Our analysis provides further evidence of this, and further suggests that the impact of the reform on the speech community may be limited in such cases.

In summary, this analysis indicates that the BwS method can be used successfully to characterise evolutionary forces that change over time from time series data alone. As an unsupervised method, it does not rely on any prior knowledge as to when the change may have occurred, although it does benefit from a large sample size being available, obtained here by aggregating multiple instances of a change together. We have found that the estimated time at which the selection strength changed corresponds well with the time at which the corresponding reform was introduced, and comparing these strengths before and after the reform allows us to assess its impact on the speech community.

4 Discussion

In this work we have applied an algorithm for the quantitative study of evolutionary time series (Montero and Blythe 2023) to instances of competition in language change. This algorithm is based on likelihood-maximisation methods and the Beta-with-Spikes (BwS) approximation to the Wright–Fisher model. The applicability of the Wright–Fisher model was justified through both theoretical considerations (Croft 2000; Hull 2010) and its manifestation as an agent-based model of language change from various starting points (Baxter et al. 2006; Blythe and Croft 2021; Reali and Griffiths 2010).

In Section 2.3, we demonstrated that the BwS method better captured the statistical properties of the Wright–Fisher model than the normal approximation that has been used elsewhere (Newberry et al. 2017). In particular, it deals better with

situations where variant frequencies are close to 0 or 1, which arises in the case when selection serves to eliminate linguistic variation across the speech community. Through refinements to the original BwS method of Tataru et al. (2015) that are detailed by Montero and Blythe (2023), we further gain accuracy in regimes where the selection strength is large.

Our first application was to the set of 36 COHA verbs previously investigated by other methods (Karjus et al. 2020; Karsdorp et al. 2020; Newberry et al. 2017). In particular, we found that even when the Frequency Increment Test (FIT, Newberry et al. 2017; Karjus et al. 2020) delivered unreliable results due to shortcomings of the normal approximation that it relies on, we obtained evidence of selection that was broadly consistent with that obtained within a Time Series Classification (TSC, Karsdorp et al. 2020) which took the complementary approach of training neural networks with artificial time series. The present method further delivered graded measure of the extent to which the historical changes are consistent with drift (in the form of a null hypothesis p -value) along with maximum likelihood estimates of parameters in the Wright–Fisher model.

A degree of care is needed when interpreting this p -value. All evolutionary trajectories are likely to be the product of some combination of drift and selection. The key question is whether their respective contributions can be distinguished. For example, a variant could be strongly selected for (large s) but subject to sufficiently large fluctuations (small N) that the systematic effects of selection are masked. The p -value is therefore a measure of the extent to which fluctuations alone could account for the changes that have been observed. If one chooses to apply the conventional significance threshold for rejecting this null hypothesis ($p < 0.05$), we find consistency with Newberry et al.'s (2017) observation that the evolution of many verbs appears to be dominated by drift.

A second important question is whether these fluctuations are a consequence of the finite number of tokens available for analysis in historical corpora, or an intrinsic property of the language dynamics within the speech community. One way to gain an insights into this question is to compare results obtained with different temporal binnings (Figures 4 and 5), since wider bins contain more points and should reduce fluctuations due to sampling. If sampling effects were dominating, we would expect to see the p -value for the drift hypothesis to decrease as the bins are widened (i.e., increasing darkness in Figure 4). This happens for some, but not all the verbs in the intermediate region, suggesting that drift may be the dominant factor in the evolution of a substantial fraction of the COHA verbs (again, consistent with Newberry et al. 2017). A more rigorous answer to this question could be obtained by incorporating finite sample-size effects into the data likelihood function used in the analysis. This is however likely to be computationally demanding, and we leave this possibility for future work.

In this work, we found plotting ellipses that indicate the variation in estimates of the selection strength and the p -value for the drift hypothesis helpful to understand which variants are more likely to have been selected for. A comparison between a baseline set of verbs from the Google Books corpus and a set where the past tense is formed by deletion of a repeated consonant reveals that they are distributed differently across the space of selection strengths s and drift p -values. Specifically, we found that the phonological simplicity arising from coalescence or omission of the /*id*/ termination tended to be favoured over the inflectional simplicity of the regular form. In principle, the method we have set out here could be used to determine the relative importance of other pairs of constraints that correspond to opposing selective forces.

Finally, we showed that the method could be applied also to changes that do not have a cognitive origin and manifest as the selection strength changing over time. We studied the dynamics of word spellings in Spanish before and after reforms introduced by their central regulatory institution, the Real Academia Española. We found that each of the changes was much better described by a model in which the selection strength changed at one or more points in time, and that the primary change point corresponded well with that at which the reform was introduced. This is despite the presence of noise on the time series data. Since changes in selection strength could derive from a variety of social and cultural factors, and indeed apply to cultural evolutionary processes beyond language, this method for automated detection of societal trends and shifts could have broad applicability.

Despite these promising results, there are inevitably some limitations. Chief among these is an inability to separate different contributions to the selective pressures acting on the system. Therefore, although it is possible to use this data to determine that selection has favoured one variant over another over time, and to estimate the strength of the effect, we have had to appeal to additional information to relate to likely causes of selection. This, however, is a problem intrinsic to the Wright–Fisher model with selection and not specific to the BwS method: the Wright–Fisher model contains only a single parameter s that characterises all systematic contributions to changes in variant frequencies.

This oversimplification of the contributing factors to language change stems, at least in part, from the underlying assumption that the competition between forms (e.g. irregular and regular verbal forms) occurs in isolation, uninfluenced by the competition dynamics of related forms (e.g. irregular and regular forms of other verbs). Yeh et al. (2019) and Buskell et al. (2019) argue in the context of cultural evolution, that cultural change may arise as an emergent phenomenon when cultural traits are interconnected. It is possible, then, that emergent system-level effects may account for significant changes in usage frequency of variants that are not favoured individually by any social or inductive bias. More refined models, ones that

account for the complex web of interconnected forms and functions present in language, may be able to differentiate between these systemic effects and those affecting individual variants. Such models might allow more information to be extracted from corpora without the need for additional information.

Nevertheless, we have shown that it is possible to draw inferences about contributions to selection from different sources (as was done in the analysis of competition between regular and irregular forms in English verbs) and quantify the impact of social factors (as was done in the language reform example). By appealing to a wider range of corpora and instances of change, it may become possible to identify general mechanisms that are invariant over time and operate cross-linguistically, and are thus informative about language universals in general. Furthermore, the method is not specific to linguistic variation, and could be used to address similar questions in other instances of cultural evolution.

Data availability

The code and data are available here: [link](#)

Acknowledgments: We thank Johanna Basnak (University of Edinburgh) for helpful discussions.

Research funding: Juan Guerrero Montero holds a Principal's Career Development Scholarship awarded by the University of Edinburgh. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Competing interests: The authors declare no competing interests.

Appendix

A Maximum likelihood parameters for the COHA verbs

In the following tables we quote the maximum-likelihood estimates of the parameters in the Wright–Fisher model obtained by applying the Beta-with-Spikes method outlined in the main text to frequency counts derived from the COHA corpus. Each table corresponds to a different binning strategy: for example, in the first table, frequency counts from each period of 10 consecutive years are aggregated to form a single frequency estimate for the corresponding time period.

Two different effective population sizes N are quoted: one ('for drift') under the assumption that $s = 0$, and the other ('for selection') that is obtained when both N and s are optimised via the maximum likelihood analysis. The p -value is the empirical p -value for the drift hypothesis, obtained as described in Section 2 of the main text. The maximum likelihood values are all quoted to three significant figures, and the p -values to two significant figures.

10-year bins

Verb	N for drift	N for selection	s	p -Value
Awake	1,820	1,990	0.021	0.11
Build	2,820	3,130	0.026	0
Burn	410	542	-0.05	0
Catch	3,690	4,170	0.031	0.028
Dive	145	147	0.0092	0.43
Draw	2,880	2,880	0.00067	0.96
Dream	219	233	-0.036	0.002
Dwell	568	554	-0.029	0.002
Grow	4,510	4,770	0.031	0.032
Hang	1,520	1,830	0.048	0
Hear	6,160	6,920	0.047	0.044
Heave	147	145	0.0069	0.68
Kneel	360	362	0.0028	0.82
Knit	121	122	-0.0043	0.77
Know	4,240	4,350	0.0035	0.65
Lay	4,070	4,250	0.002	0.47
Lean	753	926	-0.053	0
Leap	313	346	0.022	0.16
Learn	652	847	-0.052	0
Light	294	368	0.021	0.03
Plead	1,500	1,590	-0.012	0.19
Quit	790	927	0.022	0
Shine	1,330	1,320	-0.014	0.17
Smell	319	399	-0.036	0.004
Sneak	415	453	0.023	0.068
Speed	299	327	0.024	0.052
Spell	288	307	-0.017	0.31
Spill	304	357	-0.032	0
Spoil	223	226	-0.0036	0.79
Strew	127	127	-0.0023	0.86
Tell	3,760	3,960	0.018	0.37
Throw	2,090	2,180	0.012	0.25
Wake	815	928	0.015	0.012
Weave	460	457	-0.007	0.45

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Wed	116	120	0.026	0.03
Wet	201	201	0.0024	0.88

20-year bins

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Awake	5,130	5,730	0.011	0.13
Build	4,850	5,290	0.0064	0.27
Burn	615	909	-0.042	0
Catch	9,890	12,000	0.02	0.058
Dive	483	546	0.0095	0.24
Draw	6,690	6,660	0.0028	0.84
Dream	296	326	-0.034	0.004
Dwell	1,080	1,430	-0.04	0
Grow	17,600	17,600	-4.2e-05	1.0
Hang	4,260	4,310	0.0052	0.58
Hear	16,400	17,500	0.013	0.39
Heave	353	350	0.0018	0.85
Kneel	1,580	1,640	0.0018	0.76
Knit	244	251	-0.0049	0.66
Know	8,890	8,900	4.1e-05	1.0
Lay	9,110	10,200	0.0016	0.43
Lean	1,060	1,990	-0.063	0
Leap	498	616	0.019	0.1
Learn	937	1,910	-0.054	0
Light	360	834	0.022	0.006
Plead	2,100	2,110	-0.00072	0.94
Quit	1,140	1,690	0.025	0
Shine	2,500	2,740	-0.018	0.066
Smell	500	1,440	-0.036	0
Sneak	674	871	0.024	0.02
Speed	807	941	0.014	0.14
Spell	498	1,310	-0.042	0
Spill	688	1,220	-0.025	0
Spoil	565	747	-0.028	0.018
Strew	334	334	0.0012	0.87
Tell	7,620	7,570	0.012	0.62
Throw	9,380	9,420	0.00046	0.94
Wake	1,160	1,650	0.015	0.014
Weave	1,020	1,000	-0.011	0.14

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Wed	147	216	0.028	0.1
Wet	1,100	1,140	0.0021	0.78

40-year bins

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Awake	9,510	20,300	0.016	0.042
Build	5,780	8,090	0.0098	0.16
Burn	1,770	15,100	-0.017	0.014
Catch	15,000	40,800	0.027	0.002
Dive	697	879	0.0072	0.49
Draw	30,800	48,600	0.0057	0.52
Dream	824	868	-0.0075	0.52
Dwell	1,430	2,060	-0.037	0
Grow	31,200	31,100	-0.0012	0.95
Hang	146,000	148,000	-0.00099	0.68
Hear	20,100	20,300	0.0073	0.68
Heave	743	835	-0.0046	0.66
Kneel	3,090	6,490	0.0062	0.25
Knit	252	261	-0.0027	0.83
Know	8,310	8,720	0.0025	0.75
Lay	9,020	13,400	0.0021	0.35
Lean	2,810	5,530	-0.025	0.008
Leap	775	968	0.015	0.14
Learn	2,680	9,430	-0.023	0.006
Light	318	1,110	0.019	0.03
Plead	3,770	3,690	0.0047	0.62
Quit	659	1,610	0.026	0
Shine	3,910	8,610	-0.019	0.01
Smell	477	2,600	-0.034	0
Sneak	924	920	0.014	0.24
Speed	1,050	1,040	-0.00025	1.0
Spell	509	2,180	-0.039	0
Spill	712	2,060	-0.026	0
Spoil	1,350	2,490	-0.014	0.086
Strew	569	623	0.0043	0.65
Tell	11,000	11,700	0.025	0.46
Throw	6,510	6,630	-0.0017	0.93
Wake	908	2,810	0.019	0.014
Weave	1,160	1,690	-0.017	0.052

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Wed	248	364	0.016	0.14
Wet	2,570	3,220	0.0037	0.13

B Maximum likelihood parameters for verbs in the study of competing motivations

In this Appendix, we provide the corresponding tables for the set of verbs ending in alveolar stops from drawn from the Google Books corpus. Dashes mean that the corresponding time series did not have enough data points per time bin in the corresponding binning for it to be included in the study.

5-year bins

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Bend	7,450	9,980	0.019	0.004
Bet	–	–	–	–
Bite	4,600	6,960	0.029	0
Blend	5,110	5,040	0.0037	0.56
Build	21,000	22,200	0.0058	0.28
Fit	586	596	0.026	0.036
Glide	6,520	6,520	–0.0026	0.84
Knit	–	–	–	–
Light	823	908	0.011	0.036
Pat	1,420	1,420	–0.02	0.12
Plead	5,040	5,160	0.01	0.054
Quit	376	377	0.076	0
Slide	2,050	2,210	0.023	0.002
Speed	628	626	–0.003	0.61
Spit	757	847	0.014	0.064
Thrust	2,510	2,880	0.074	0.002
Tread	2,980	3,000	–0.014	0.25
Wed	93.6	83.6	0.079	0
Wet	–	–	–	–
Awake	1,390	2,230	0.025	0
Blow	11,100	11,200	0.0012	0.8
Burn	1,110	1,350	–0.012	0
Catch	5,070	6,640	0.039	0
Cleave	442	447	–0.0042	0.54
Creep	3,820	4,190	0.042	0.014

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Dive	757	766	0.011	0.088
Dream	1,830	1,900	-0.0041	0.36
Dwell	1,300	1,300	0	1.0
Freeze	2,290	2,510	0.045	0.008
Grow	95,400	95,300	-0.0012	0.5
Hang	2,070	2,350	0.0077	0.052
Heave	507	516	-0.0042	0.58
Hew	1,320	1,320	-0.00037	0.93
Kneel	986	1,390	0.02	0.002
Lean	1,350	1,380	-0.0037	0.47
Leap	1,460	1,500	0.0056	0.22
Learn	2,250	2,360	-0.0047	0.3
Shake	3,910	4,520	0.065	0
Shear	545	546	-0.0073	0.26
Shine	1,380	1,380	0.00025	0.98
Slay	8,430	8,430	-0.00048	0.98
Slink	1,310	1,320	-0.011	0.28
Smell	710	799	-0.013	0.018
Sneak	2,050	3,130	0.055	0
Spell	542	584	-0.011	0.1
Spill	863	1,260	-0.02	0.002
Spoil	1,370	1,380	0.0046	0.26
Strew	646	1,010	0.028	0
String	2,180	2,650	0.019	0.038
Strive	3,520	3,830	-0.017	0.026
Swell	1,070	1,240	0.011	0.01
Wake	775	1,290	0.025	0
Weave	994	988	-0.0086	0.16

10-year bins

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Bend	12,400	28,100	0.017	0
Bet	491	685	0.039	0
Bite	5,330	16,600	0.03	0
Blend	8,290	8,170	0.0024	0.68
Build	23,900	27,000	0.0057	0.23
Fit	1,340	1,750	0.03	0
Glide	26,300	25,900	0.0014	0.85
Knit	-	-	-	-

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Light	840	1,120	0.012	0.018
Pat	2,680	3,090	-0.022	0.028
Plead	6,390	7,030	0.011	0.028
Quit	522	670	0.052	0
Slide	3,240	3,240	0.018	0.012
Speed	444	445	-0.0024	0.75
Spit	1,210	1,590	0.013	0.054
Thrust	3,760	4,630	0.051	0.006
Tread	7,190	8,840	-0.023	0
Wed	159	163	0.038	0.038
Wet	-	-	-	-
Awake	1,310	3,710	0.025	0
Blow	24,500	24,700	0.00079	0.81
Burn	1,260	2,260	-0.013	0
Catch	9,740	28,500	0.033	0
Cleave	461	472	-0.0042	0.51
Creep	12,600	14,200	0.017	0.066
Dive	1,050	1,140	0.012	0.048
Dream	2,650	3,100	-0.0053	0.16
Dwell	1,790	1,970	-0.014	0.096
Freeze	6,130	7,090	0.02	0.038
Grow	184,000	187,000	-0.0018	0.19
Hang	3,010	4,150	0.0071	0.04
Heave	1,130	1,320	-0.007	0.14
Hew	7,420	7,640	-0.0015	0.43
Kneel	895	1,560	0.019	0.008
Lean	1,350	1,420	-0.0038	0.49
Leap	1,660	1,830	0.0066	0.1
Learn	3,130	3,660	-0.0054	0.17
Shake	9,620	11,900	0.037	0.004
Shear	983	1,100	-0.0098	0.06
Shine	3,570	3,550	-0.00087	0.9
Slay	17,200	20,000	-0.015	0.21
Slink	2,120	2,230	-0.015	0.12
Smell	607	785	-0.014	0.022
Sneak	2,190	3,740	0.055	0
Spell	998	1,630	-0.014	0.004
Spill	789	1,890	-0.02	0
Spoil	1,710	1,730	0.0039	0.26
Strew	453	1,130	0.03	0
String	8,410	12,200	0.012	0.028
Strive	2,980	3,440	-0.018	0.034
Swell	844	1,070	0.011	0.03

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Wake	631	1,710	0.025	0
Weave	2,110	2,280	-0.0093	0.044

20-year bins

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Bend	13,800	126,000	0.016	0
Bet	677	913	0.02	0.056
Bite	5,670	18,200	0.033	0
Blend	5,580	5,580	0.0013	0.88
Build	22,600	28,500	0.0064	0.26
Fit	1,390	2,490	0.032	0
Glide	102,000	116,000	-0.0035	0.58
Knit	1,460	1,480	-0.0043	0.72
Light	746	1,850	0.014	0.008
Pat	2,660	6,670	-0.039	0
Plead	5,860	7,540	0.011	0.044
Quit	590	1,020	0.054	0
Slide	5,160	23,200	0.026	0
Speed	434	430	-0.0042	0.63
Spit	1,570	9,000	0.017	0
Thrust	4,700	4,250	-0.013	0.52
Tread	8,800	21,600	-0.028	0
Wed	310	353	0.037	0.012
Wet	800	898	0.0051	0.6
Awake	1,150	3,460	0.023	0.006
Blow	26,200	26,700	0.00064	0.87
Burn	1,040	3,790	-0.013	0.004
Catch	12,500	47,200	0.03	0
Cleave	374	399	-0.0074	0.41
Creep	11,200	13,400	0.016	0.14
Dive	1,120	1,550	0.012	0.046
Dream	3,030	5,060	-0.0061	0.074
Dwell	2,570	4,140	-0.02	0.01
Freeze	14,800	24,300	0.013	0.086
Grow	190,000	190,000	-0.0022	0.36
Hang	3,850	8,570	0.0064	0.022
Heave	1,530	2,210	-0.0064	0.12
Hew	9,390	10,200	-0.0016	0.4
Kneel	978	2,630	0.015	0.02

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Lean	1,840	2,300	-0.0059	0.28
Leap	3,010	5,730	0.007	0.02
Learn	5,100	8,270	-0.0052	0.072
Shake	22,800	32,600	0.021	0.026
Shear	1,550	2,980	-0.011	0.008
Shine	2,570	2,540	-0.002	0.78
Slay	12,100	162,000	-0.026	0.01
Slink	3,280	5,200	-0.027	0.004
Smell	487	863	-0.015	0.018
Sneak	2,750	4,340	0.056	0
Spell	890	3,350	-0.014	0.004
Spill	624	7,300	-0.02	0
Spoil	1,860	1,910	0.0035	0.43
Strew	295	1,110	0.031	0.006
String	12,100	20,400	0.0086	0.1
Strive	3,730	4,800	-0.021	0.008
Swell	618	828	0.01	0.1
Wake	476	1,350	0.025	0
Weave	2,450	2,850	-0.0081	0.12

40-year bins

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i> -Value
Bend	12,900	196,000	0.014	0.008
Bet	748	3,900	0.025	0.008
Bite	7,370	38,000	0.034	0
Blend	5,020	5,030	-0.00054	0.98
Build	18,600	26,800	0.0082	0.27
Fit	1,880	2,450	0.027	0
Glide	70,600	77,800	-0.0024	0.74
Knit	2,120	10,400	-0.017	0.13
Light	366	2,180	0.016	0.014
Pat	3,470	67,700	-0.03	0
Plead	6,110	16,700	0.0099	0.054
Quit	966	1,460	0.049	0
Sit	200	200	0.17	0.99
Slide	5,730	36,500	0.028	0
Speed	487	443	-0.014	0.21
Spit	1,580	27,000	0.017	0.004
Thrust	17,100	26,400	0.054	0.008

(continued)

Verb	<i>N</i> for drift	<i>N</i> for selection	<i>s</i>	<i>p</i>-Value
Tread	9,620	26,800	-0.033	0
Wed	513	702	0.035	0.002
Wet	1,340	1,620	0.0039	0.16
Awake	1,810	13,500	0.017	0.022
Blow	14,500	15,200	0.00087	0.9
Burn	758	5,850	-0.012	0.024
Catch	18,500	101,000	0.023	0.004
Cleave	353	408	-0.0077	0.52
Creep	16,100	16,300	0.00068	0.98
Dive	700	1,130	0.012	0.12
Dream	2,780	5,950	-0.0052	0.17
Dwell	2,670	3,730	-0.022	0.018
Freeze	20,200	86,700	0.014	0.058
Grow	112,000	110,000	-0.0028	0.5
Hang	3,720	28,400	0.0059	0.028
Heave	1,210	4,750	-0.0083	0.076
hew	8,230	11,300	-0.0021	0.31
Kneel	965	3,050	0.012	0.16
Lean	1,640	3,080	-0.0082	0.21
Leap	2,410	20,700	0.0074	0.012
Learn	4,610	10,100	-0.0046	0.16
Shake	24,700	69,000	0.025	0.006
Shear	727	8,120	-0.013	0.01
Shine	1,920	1,920	-0.00042	0.98
Slay	16,000	461,000	-0.028	0.014
Slink	3,860	4,260	-0.026	0.004
Smell	290	599	-0.016	0.1
Sneak	3,790	4,200	0.059	0.004
Spell	497	2,340	-0.014	0.044
Spill	385	10,600	-0.019	0
Spoil	1,260	1,230	0.0034	0.58
Strew	300	656	0.03	0.016
String	25,700	59,300	0.0053	0.22
Strive	3,990	4,020	-0.022	0.022
Swell	565	804	0.0081	0.36
Wake	506	2,030	0.022	0.008
Weave	1,350	1,830	-0.0087	0.27

C Sets of words used in RAE reform detections

Old spellings of words in set (A): <ss>to <s>

*assegurar, assentar, assentir, assunto, confessar, diesse, essa, esse, essencia, esso, estu-
viessse, fuesse, gustasse, hiciesse, passar, pudiesse, quisiesse, tassar, tuviessse, and usasse.*

Old spellings of words in set (B): <x>to <j>

*abaxo, baxar, baxo, bruxa, bruxería, caxa, conduxo, debaxo, dexar, dibuxar, dibuxo,
dixo, enxuto, exe, exemplo, exercer, ejercicio, ejército, floxo, fluxo, fixar, fixo, quexa,
roxo, texa, and traxo.*

Old spellings of words in set (C): <y>to <i>

*aceyte, aceytuna, afeyte, amaynar, ayre, bayle, deleyte, deydad, estoyco, frayle, gayta,
heroyco, layco, oyga, peyne, and reyna.*

Old spellings of words in set (D.1): Word-final tonic syllable becoming accentuated

*accion, alacran, algun, almacen, atencion, bailarín, cancion, capitán, comun, corazón,
estacion, jardín, latin, nacion, ningun, opcion razon, recien, region, relacion, segun,
Serafín, situacion, tambien, and union.*

Old spellings of words in set (D.2): Non word-final tonic syllable losing its accent

*abdómen, álguien, Cármen, certámen, cólon, crímen, desórden, dictámen, exámen,
gérmen, jóven, márgen, órden, orígen, resúmen, and volúmen.*

D Maximum likelihood parameters for time series in the study of Spanish spelling reforms

Time series	<i>T</i>	<i>N</i> before	<i>s</i> before	<i>N</i> after	<i>s</i> after	<i>p</i> -Value
(A)	1,775	5.0	-0.008	224	-0.37	<0.002
	1,805	60	-0.59	292	-0.10	0.024
(B)	1,825	14	-0.29	169	-0.072	<0.002
	1,810	17	0.003	33	-2.2	0.036
	1,840	2,530	-1.0	154	-0.02	0.026
(C)	1,825	9.3	-0.30	121	-0.061	<0.002
(D1)	1,890	16	-0.050	427	-0.21	<0.002
(D2)	1,890	85	0.13	249	-0.53	<0.002

References

- Acerbi, Alberto & Alexander Bentley. 2014. Biases in cultural transmission shape the turnover of popular traits. *Evolution and Human Behavior* 35(3). 228–236.
- Amato, Roberta, Lucas Lacasa, Albert Díaz-Guilera & Andrea Baronchelli. 2018. The dynamics of norm change in the cultural evolution of language. *Proceedings of the National Academy of Sciences of the United States of America* 115(33). 8260–8265.
- Anderwald, Lieselotte. 2012. Variable past-tense forms in nineteenth-century American English: Linking normative grammars and language change. *American Speech* 87(3). 257–293.
- Atkinson, Quentin D. & Russell D. Gray. 2005. Curious parallels and curious connections – phylogenetic thinking in biology and historical linguistics. *Systematic Biology* 54(4). 513–526.
- Baddeley, Susan & Anja Voeste. 2012. *Orthographies in Early Modern Europe*. De Gruyter Mouton.
- Bates, Elizabeth & Brian MacWhinney. 1987. Competition, variation and language learning. In Brian MacWhinney (ed.), *Mechanisms of language acquisition*, 157–193. Hillsdale, NJ: Lawrence Erlbaum.
- Bates, Elizabeth & Brian MacWhinney. 1989. Functionalism and the competition model. In Brian MacWhinney & Elizabeth Bates (eds.), *The crosslinguistic study of sentence processing*, 3–73. Cambridge: Cambridge University Press.
- Baxter, Gareth J., Richard A. Blythe, William Croft & Alan J. McKane. 2006. Utterance selection model of language change. *Physical Review E* 73. 046118.
- Blythe, Richard A. 2012. Neutral evolution: A null model for language dynamics. *Advances in Complex Systems* 15(3–4). 1150015.
- Blythe, Richard A. & William Croft. 2021. How individuals change language. *PLoS One* 16(6). 1–23.
- Boyd, Robert & Peter J. Richerson. 1988. *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Bromham, Lindell, Xia Hua, Thomas G. Fitzpatrick & Simon J. Greenhill. 2015. Rate of language evolution is affected by population size. *PNAS* 112. 2097–2102.
- Buskell, Andrew, Magnus Enquist & Fredrik Jansson. 2019. A systems approach to cultural evolution. *Palgrave Communications* 5(1). 131.
- Bybee, Joan. 1995. Regular morphology and the lexicon. *Language and Cognitive Processes* 10. 425–455.
- Bybee, Joan. 2001. *Phonology and language use*. Cambridge: Cambridge University Press.
- Bybee, Joan. 2007. *Frequency of use and the organization of language*. Oxford: Oxford University Press.
- Cavalli-Sforza, Luigi Luca & Marcus W. Feldman. 1981. *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.
- Croft, William. 2000. *Explaining language change: An evolutionary approach*. London: Pearson Education.
- Crow, James F. & Motoo Kimura. 1970. *An introduction in population genetics theory*. New York: Harper & Row.
- Cuskley, Christine F., Martina Pugliese, Claudio Castellano, Francesca Colaiori, Vittorio Loreto & Francesca Tria. 2014. Internal and external dynamics in language: Evidence from verb regularity in a historical corpus of English. *PLoS One* 9(8). e102882.
- Davies, Mark. 2010. *The corpus of historical American English*. Available at: <https://www.englishcorpora.org/coha/>.
- DuBois, John W. 1985. Competing motivations. In John Haiman (ed.), *Iconicity in syntax*, 343–366. Amsterdam: John Benjamins.
- Feder, Alison F., Sergey Kryazhimskiy & Joshua B. Plotkin. 2014. Identifying signatures of selection in genetic time series. *Genetics* 196(2). 509–522.

- Fisher, Ronald A. 1930. *The genetical theory of natural selection*. Oxford: Clarendon Press.
- Frisch, Stephan A., Janet B. Pierrehumbert & Michael B. Broe. 2004. Similarity avoidance and the ocp. *Natural Language & Linguistic Theory* 22. 179–228.
- Hahn, Matthew W. & Alexander Bentley. 2003. Drift as a mechanism for cultural change: An example from baby names. *Proceedings of the Royal Society of London B: Biological Sciences* 270. S120–S123.
- Haiman, John. 1983. Iconic and economic motivation. *Language* 53. 781–819.
- Hawkins, John A. 2004. *Efficiency and complexity in grammars*. Oxford: OUP Oxford.
- Hayes, Bruce P. 1999. Phonetically driven phonology. *Functionalism and Formalism in Linguistics* 1. 243–285.
- Hernández-Campoy, Juan Manuel & Juan Camilo Conde-Silvestre. 2012. *The handbook of historical sociolinguistics*, 68. Hoboken: John Wiley & Sons.
- Hull, David L. 2010. *Science as a process: An evolutionary account of the social and conceptual development of science*. Chicago: University of Chicago Press.
- Jolliffe, Ian T. 2002. *Principal component analysis*. New York: Springer.
- Kandler, Anne & Adam Powell. 2015. Inferring learning strategies from cultural frequency data. In Alex Mesoudi & Kenichi Aoki (eds.), *Learning strategies and cultural evolution during the Palaeolithic*, 85–101. Tokyo: Springer Japan.
- Kandler, Anne & Stephen Shennan. 2013. A non-equilibrium neutral model for analysing cultural change. *Journal of Theoretical Biology* 330. 18–25.
- Karjus, Andres, Richard A. Blythe, Simon Kirby & Kenny Smith. 2020. Challenges in detecting evolutionary forces in language change using diachronic corpora. *Glossa* 5. 45.
- Karsdorf, Folgert, Enrique Manjavacas, Lauren Fonteyn & Mike Kestemont. 2020. Classifying evolutionary forces in language change using neural networks. *Evolutionary Human Sciences* 2. E50.
- Kauhanen, Henri & George Walkden. 2018. Deriving the constant rate effect. *Natural Language & Linguistic Theory* 36. 483–521.
- Keller, Rudi & Brigitte Nerlich. 1994. *On language change: The invisible hand in language*. Abingdon-on-Thames: Routledge.
- Kirby, Simon. 1997. Competing motivations and emergence: Explaining implicational hierarchies. *Linguistic Typology* 1. 995–1026.
- Kroch, Anthony S. 1989. Reflexes of grammar in patterns of language change. *Language Variation and Change* 1(3). 199–244.
- Labov, William. 1994. *Principles of linguistic change*, vol. 1. Oxford: Blackwell.
- Labov, William. 2001. *Principles of linguistic change*, vol. 2. Oxford: Blackwell.
- Labov, William. 2010. *Principles of linguistic change*, vol. 3. Oxford: Blackwell.
- Leben, William R. 1973. *Suprasegmental phonology*. Cambridge: Massachusetts Institute of Technology dissertation.
- Lieberman, Erez, Jean-Baptiste Michel, Joe Jackson, Tina Tang & Martin A. Nowak. 2007. Quantifying the evolutionary dynamics of language. *Nature* 449(7163). 713–716.
- McCarthy, John J. 1986. Ocp effects: Gemination and antigemination. *Linguistic Inquiry* 17. 207–264.
- McDonald, John H. 2014. *Handbook of biological statistics*. Baltimore: Sparky House Publishing. Chap. G-test of goodness-of-fit.
- McMahon, April M. S. 1994. *Understanding language change*. Cambridge: Cambridge University Press.
- Michel, Jean-Baptiste, Yuan K. Shen, Aviva P. Aiden, Adrian Veres, Matthew K. Gray, Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak & Erez L. Aiden. 2011. Quantitative analysis of culture using millions of digitized books. *Science* 331(6014). 176–182.
- Montero, Juan Guerrero & Richard A. Blythe. 2023. Self-contained Beta-with-Spikes approximation for inference under a Wright–Fisher model. *Genetics* 225. <https://doi.org/10.1093/genetics/iyad092>.

- Mufwene, Salikoko S. 2001. *The ecology of language evolution*. Cambridge: Cambridge University Press.
- Newberry, Mitchell, Christopher Ahern, Robin Clark & Joshua B. Plotkin. 2017. Detecting evolutionary forces in language change. *Nature* 551. 223–226.
- Pagel, Mark. 2009. Human language as a culturally transmitted replicator. *Nature Reviews Genetics* 10(6). 405–415.
- Paris, Cyriel, Bertrand Servin & Simon Boitard. 2019. Inference of selection from genetic time series using various parametric approximations to the Wright–Fisher model. *G3 Genes—Genomes—Genetics* 9(12). 4073–4086.
- Pechenick, Eitan Adam, Christopher M. Danforth & Peter Sheridan Dodds. 2015. Characterizing the Google Books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PLoS One* 10(10). 1–24.
- Pozdniakov, Konstantin & Guillaume Segerer. 2007. Similar place avoidance: A statistical universal.
- Prasada, Sandeep & Steven Pinker. 1993. Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* 8. 1–56.
- Prince, Alan & Paul Smolensky. 1997. Optimality: From neural networks to universal grammar. *Science* 275(5306). 1604–1610.
- Real Academia Española. 1763. *Ortografía de la lengua castellana*, 3rd edn. Madrid, Spain: Imprenta Real.
- Real Academia Española. 1815. *Ortografía de la lengua castellana*, 8th edn. Madrid, Spain: Imprenta Real.
- Real Academia Española. 1881. *Prontuario de ortografía castellana en preguntas y respuestas*, 7th edn. Madrid: Gregorio Hernando.
- Realí, Florencia & Thomas L. Griffiths. 2010. Words as alleles: Connecting language evolution with Bayesian learners to models of genetic drift. *Proceedings of the Royal Society B* 277. 429–436.
- Ringe, Don & Charles Yang. 2022. The threshold of productivity and the ‘irregularization’ of verbs in Early Modern English. In Bettelou Los, Claire Cowie, Patrick Honeybone & Graeme Trousdale (eds.), *English historical linguistics: Change in structure and meaning*. Amsterdam: John Benjamins.
- Rubin, Joan, Björn H. Jernudd, Jyotirindra DasGupta, Joshua A. Fishman & Charles A. Ferguson (eds.). 2013. *Language planning processes*. Berlin: De Gruyter Mouton.
- Rutten, Gijsbert & Rik Vosters. 2021. Language standardization ‘from above’. *The Cambridge handbook of language standardization*, 65–92. Cambridge: Cambridge University Press.
- Sapir, Edward. 1921. *Language: An introduction to the study of speech*. San Diego: Harcourt.
- Severini, Thomas A. 2000. *Likelihood methods in statistics*. New York: Oxford University Press.
- Silvey, Samuel D. 1970. *Statistical inference*. London: Chapman & Hall.
- Sims-Williams, Helen. 2016. Analogical levelling and optimisation: The treatment of pointless lexical allomorphy in Greek. *Transactions of the Philological Society* 114(3). 315–338.
- Steele, James, Claudia Glatz & Anne Kandler. 2010. Ceramic diversity, random copying, and tests for selectivity in ceramic production. *Journal of Archaeological Science* 37(6). 1348–1358.
- Stemberger, Joseph Paul. 1981. Morphological haplology. *Language* 57(4). 791–817.
- Tagliamonte, Sali A. 2011. *Variationist sociolinguistics: Change, observation, interpretation*. Hoboken: Wiley.
- Tataru, Paula, Thomas Bataillon & Asger Hobolth. 2015. Inference under a Wright–Fisher model using an accurate Beta approximation. *Genetics* 201. 1133–1151.
- Tataru, Paula, Maria Simonsen, Thomas Bataillon & Asger Hobolth. 2016. Statistical inference in the Wright–Fisher model using allele frequency data. *Systematic Biology* 66(1). e30–e46.
- Taylor, Wayne A. 2000. *Change-point analysis: A powerful new tool for detecting changes*. Libertyville: Taylor Enterprises.
- Walker, James A. 2010. *Variation in linguistic systems*, 1st edn. Abingdon-on-Thames: Routledge.
- Wichmann, Soeren, Dietrich Stauffer, Christian Schulze & Eric W. Holman. 2008. Do language change rates depend on population size? *Advances in Complex Systems* 11. 357–369.

- Wright, Sewall. 1931. Evolution in Mendelian populations. *Genetics* 16(2). 97–159.
- Yang, Charles. 2000. Internal and external forces in language change. *Language Variation and Change* 12(3). 231–250.
- Yang, Charles. 2002. Grammar competition and language change. In David W. Lightfoot (ed.), *Syntactic effects of morphological change*, 343–366. Oxford: Oxford University Press.
- Yeh, Justin D., Laurel Fogarty & Anne Kandler. 2019. Cultural linkage: The influence of package transmission on cultural dynamics. *Proceedings of the Royal Society B* 286. 20191951.
- Zipf, George K. 1949. *Human behavior and the principle of least effort*. Boston: Addison-Wesley Press.