

Survival factors in the early Middle English lexicon

JOHANNA VOGELSANGER 

University of Zurich

(Received 28 March 2022; revised 1 February 2023)

Using a corpus linguistic approach, this article aims to answer the question of which factors contribute to a better chance of survival for words in the early Middle English lexicon. Because of the cognitive benefits of rhyme that have been shown in modern studies, there is a particular interest in rhyming position as a potential factor; other factors include frequency, suffix and geographical spread. The data are analysed using survival analysis, random forests and conditional inference trees in R. The results show that geographical spread is the most important factor, usually in combination with particular suffixes. Rhyme is not generally a significant factor in the same vein, and its importance seems to be restricted to individual cases.

Keywords: Middle English, French, lexical loss and stability, survival analysis, rhyme

1 Introduction

The upheaval in the early Middle English (ME) lexicon is a well-known phenomenon; Old English words disappear en masse, the numbers of French (and to a lesser extent Latin and Old Norse) loanwords increase rapidly, and home-grown Middle English word-formation processes are very productive (see e.g. Dekeyser & Pauwels 1990 for an overview). The commonly cited reasons for this are ‘need and prestige’ (Durkin 2020: 172). Although prestige or need alone is not enough to bring about the adoption of French loanwords if large parts of the population never have the need or opportunity to come into contact with the language, Durkin (2020: 169) points out the difference between the ‘initial interlinguistic transfer’ and the ‘subsequent intralinguistic spread’. In the case of early Middle English, the first process would not have affected the large group of monolingual English speakers, but the second one certainly would have. Consequently, bilingual mediators would have played an important role in bringing the monolingual English populace into contact with French elements, since ‘words and structures need not necessarily be borrowed directly from the native speakers of the source language’ (Skaffari 2017: 191). The role of the clergy as such bilingual mediators has recently been highlighted (Ingham 2018), as they were educated in French, or at least had to know certain religious terminology in French and were at the same time responsible for the instruction of monolingual English laypeople (Timofeeva 2020a).

While ‘need’ is a self-explanatory reason for borrowing, there are also many cases where it could not have been the main factor since Old English (OE) had perfectly serviceable words already, and ‘prestige’ must serve as an explanation. For example, Durkin (2014: 238) describes the synonym trio *frith*, *grith* and *peace*:

Of these, *frith* barely survives beyond the early Middle English period, in one or two very restricted specialised uses. *Grith* remains quite frequent in Middle English (and occurs in some very limited contexts even later than this), although in later Middle English it occurs largely in verse texts in which its use was probably stylistically marked. (Later uses of both *frith* and *grith* are often in formulae paired with a more recent synonym, e.g. *frith and grith* or *peace and grith*.)

Of the three, *frith* is the OE word, *grith* was borrowed from Old Norse (ON) early on and already well attested in OE, so it seems that even though English had two available options to express the concept ‘peace’ (and therefore no ‘need’ to borrow from French), the French loanword was adopted and ultimately spread far enough to become part of the English lexicon. Following this development, the OE synonyms slowly fell out of use, surviving longest in poetry and binomial formulae.

A similar case is that of the trio *envy*, *nith* and *onde*: one French and two OE options, of which the French ultimately prevails and the others often appear as part of a binomial (Timofeeva 2020b). Here, too, the use in poetry warrants notice, specifically the use in rhyming position: firstly, ‘[t]he perceptual salience of rhyme is a well-known fact, which, in our case, may have contributed to the lexical priming of *onde* and, in turn, its easier acquisition, wider diffusion, and, ultimately, longer life span’ (2020b: 63), and secondly, the French synonym *envy* far surpasses either OE synonym in terms of actual and potential rhyming partners (Timofeeva 2020b: 63). This seems to fit neatly with the numbers we find in the *Linguistic Atlas of Early Middle English (LAEME; Laing 2013–)*, where *envy* is rhymed in 18/31 instances (58%), *onde* 24/71 (34%) and *nith* 10/53 (19%).

However, the replacement of Germanic with Romance vocabulary as in these examples is neither complete nor systematic. While prestige may serve as an explanation for why *envy* was borrowed, it does not explain why other, similar words in the same semantic domain were not also replaced by French borrowings. For many of the French loans in these domains, we find English near-synonyms which then go on to co-exist, and in some cases we can observe the continued survival of a word derived from Old English, despite supposed competition¹ from a French loan, even in such domains that are supposedly heavily dominated by French. For example, among the cardinal sins, which can be put very firmly into the semantic domain of ‘religion’, we find surviving English-origin *wrath* and *sloth*, which had French competitors in the words *erour* and *accyde*, respectively (Käsmann 1961: 296–300), but also French *gluttony* and *envie*, which eventually ousted their English-origin counterparts, *giferresse* and *nip* or *onde*,

¹ Characterising the changes in the ME lexicon in terms of ‘competition’ or ‘rivalry’ has recently been challenged by Sylvester *et al.* (2021).

respectively (Käsmann 1961: 293–5). What, if anything, gave one word an advantage over a contemporary (near-)synonym in cases like these?

As implied in the examples described above, it is possible that rhyme could be one such advantageous factor. In addition to these examples, there are others who mention the importance of rhyme in individual cases, e.g. Käsmann (1961: 99) on *christianity*, or Gardner (2014) with regard to competing variants of *-hood*, but so far there has not been a quantitative study investigating this possibility. Using a corpus linguistic approach, this article aims to answer the question of which factors contribute to a better chance of survival for words in the early Middle English lexicon, with a particular interest in rhyming position. I will first elaborate on the effect of rhyme on mental processing, then describe my data collection method and the statistical model used for analysis, before finally discussing the results in general as well as a few specific examples.

2 Rhyme and poetic language

Words in rhyming position are very salient, as evidenced by, for example, the long-lasting and continuous use of rhyme as a mnemonic device. Awareness of rhyme is also associated with successful language acquisition; rhyme judgement tasks are used in studies of child language acquisition to gauge their phonological awareness (e.g. Carroll & Snowling 2001), and children's ability to recognise and judge rhyme 'appears to be essentially adult-like by the age of 7' (Coch & Gullick 2011: 497 and references therein). Recent studies on the effects of rhyme on language and cognition show that, overall, rhyme lessens the cognitive load (which manifests as, e.g., shorter response times in experiments) in spoken, written and cross-modal language settings, as well as in experiments with non-words (Cutler, Van Ooijen & Norris 1999; Slowiaczek *et al.* 2000; Rapp & Samuel 2002; Coch *et al.* 2005; Mitra & Coch 2018). It should therefore be possible to investigate rhyme in medieval texts that we can only access in written form, despite our uncertainty regarding their mode of delivery and reception, and the degree to which the audience was familiar with the sound and structure of new loanwords.

As an aside, it should also be mentioned here that there are similar effects that can be observed in the case of word-initial overlap, i.e., that it facilitates recall, memorisation and attention (Slowiaczek *et al.* 2000: 533; Lea *et al.* 2008: 713; Egan *et al.* 2020). In poetry, this is present in the form of alliteration, which was of course a very important feature of Old English poetry. And while there are fewer alliterating poems in the (early) ME period than in the OE period, there still are a few, some of which even combine both alliteration and end-rhyme. For this study, however, I will only be focusing on end-rhyme.

Another aspect of poetry that is relevant in this context is its specialised vocabulary. It is well known that the Old English lexicon has a large number of words that appear only once or very rarely, and only in poetry (Scragg 2013: 60). In Middle English, too, the corpus of texts associated with the Alliterative Revival 'possesses a rich inventory of poetic lexis that is genetically derived from the specialised diction of Old English poetry' (Pascual 2017: 250). In general, if such words are part of a poetic register and/or nonce formations created to fill a particular gap in the metre or rhyme scheme, we

would not expect them to gain traction and become part of the common lexicon. Indeed, due to their limited use, these words are more likely to appear very short-lived in a study like this.²

Nevertheless, given the overall cognitive benefits of rhyme for language processing, I assume that this could have had an effect in the ME period, too. More specifically, I assume that being used in rhyme would increase a word's overall salience, which would make it easier to remember and access, and thus increase its chances of survival.

3 Data

The corpus used for this study is the electronic version of the *Linguistic Atlas of Early Middle English, 1150–1325 (LAEME)*, which contains about 650,000 words. In addition to syntactic and morphological information, the tagging in *LAEME* also includes information on whether a word appears in rhyming position. Approximately 52 per cent of all words in the corpus appear in verse texts, and 48 per cent in prose texts (the few alliterating text(s) were counted as verse for this purpose; this concerns 2,428 words in text #172 and 364 words in #136).³

Since the focus is on nouns in rhyming position, I collected my data based on nominal suffixes: *-ship*, *-ness*, *-hood*, and *-dom*, the most frequent both in terms of types and tokens (Gardner 2014: 71), for the Old and Middle English subsets (i.e. words that are present in the ME lexicon which were already attested in OE, and words which consist of Germanic elements but are for the first time attested in ME, respectively), and *-ity*, *-ery* and *-ment* for the French subset. The lexels⁴ in *LAEME* follow the modern-day spelling for words that still exist, but in those cases where a word (both of French and English origin) did not survive until today, a non-standardised spelling is used for the lexel, e.g. *vilte*. In this case, the word goes back to Latin *vilitas* (*Französisches Etymologisches Wörterbuch (FEW)*, s.v. *vilitas*, 'lowness of price or value'), but the lexel's suffix is not spelled <(i)ty>, which is the usual spelling in Present-day English for words derived from Latin words ending in *-(i)tas* (Durkin 2014: 239) such as *dignity* < *dignitas*. In order to be able to include such cases, as well as to have a subset that is comparable in size to the OE and ME ones, I expanded my criteria to include French nouns appearing in the *LAEME* lexel list that end in *-i*, *-y* or *-e*.

For each word, I also collected information about its first and last dates of attestation from the *Historical Thesaurus of English (HTE)*. Where the *HTE* and *LAEME* disagreed, I took the earlier date of the two for the first attestation and the later date of the two for the last attestation. In *LAEME*, dates are given in the following format: C13b1 stands for the first quarter (1) in the second half (b) of the thirteenth century (C13). In cases where the *LAEME* dating was used, I chose the most extreme possible

² I would like to thank an anonymous reviewer for pointing this out.

³ Calculations based on numbers given to me by M. Laing (personal communication).

⁴ A blend of *lexical* and *element*; the term describes the lexical half of a tag, as opposed to the grammatical (Laing & Lass 2007).

date for the given range. For example, the *HTE* dates *scendness* only as ‘OE’, meaning that it is only used in Old English texts, but *LAEME* has three attestations, the latest of which is dated to the first quarter of the fourteenth century (in text #286). In this case, I used 1325 as the date of last attestation.⁵

Since *LAEME* also provides a localisation of a text’s language wherever possible, I also included information about a word’s geographical spread in my dataset. The localisation is given both as a placename as well as a six-digit National Grid reference (Laing 2007) which essentially functions as a set of x- and y-coordinates whose origin point is located to the southwest of Cornwall. For example, *steadfastness* occurs in text #150, which is localised to West Norfolk with the grid reference 579 307. The first set of three numbers serve as x-coordinates and the second set as y-coordinates. When this information was given for instances in at least two different locations, I used these coordinates to determine the geographical spread of a word. To quantify this, I calculated the diagonal of a rectangle that encloses the most distant x- and y-coordinates. The advantage of using the diagonal over the area is that this more accurately represents the wide distribution of words whose coordinates on one axis happen to be very far apart while those on the other axis happen to be very close. Consider, for example, the geographical distribution of *ointment* and *steadfastness* (figures 1 and 2⁶): *ointment* is attested both in the North and in the southern Midlands (marked by red X’s), and *steadfastness* is attested in the East and West Midlands. The longer sides of the two rectangles (in green) are roughly the same length (255 km and 213 km, respectively). However, since the shorter side of the rectangle encompassing the two texts which contain *ointment* is only half the length of the shorter side of the rectangle encompassing the texts containing *steadfastness* (26 km and 52 km, respectively), its area is also only about half of the other (6,630 km² and 11,076 km², respectively). Using these numbers would give the misleading impression that *steadfastness* has twice the geographical spread of *ointment*. If we calculate the diagonal (in blue), however, the two are much more similar (256 km and 219 km, respectively). The log value of these diagonals was used for the analyses.

Finally, I recorded the absolute number of occurrences of each word in *LAEME*, as well as the number of times each word occurred in rhyming position. For the analysis, I used the log value of the absolute frequencies, as well as the rhyme-to-total ratio. One decision regarding the data collection that I had to take at some point was whether to include or discard duplicates of the kind where excerpts from the same texts from different manuscripts and/or dialect areas were included in *LAEME*. For example, *swicolhood* appears four times in *LAEME*, and is rhymed in all four cases, but only half of these

⁵ Cf. Dekeyser & Pauwels (1990: 11), who point out that, since we are dealing only with written data, not spoken, ‘lexical innovation has to be antedated, while loss has to be postdated’. I am not going so far as to modify the numbers that I found in *HTE* or *LAEME*, but I am taking the earliest possible date for first attestations and the latest possible date for last attestations.

⁶ These two figures are modified versions of *LAEME*’s Key Map (www.lel.ed.ac.uk/ihd/laeme2/MAPS/TEXTMAPS/laeme1_keymap_mon.pdf; used with permission of the copyright holder: The Angus McIntosh Centre for Historical Linguistics, The University of Edinburgh).



Figure 1. Geographical spread of *ointment*

are strictly speaking distinct because they are from the same text passage of *The Owl and the Nightingale*, which appears twice in LAEME (texts #2 and #1100). This distorts the rhyme ratio somewhat, but I felt it was necessary to keep these for the geographical spread as well as the absolute frequency counts.

The final dataset contained 118 words in the French subset, 164 in the ME subset and 193 words in the OE subset.

4 Method

In order to test whether rhyme had an effect on a word's chances of survival, I used two different approaches in R: survival analysis (SA) and random forests (RF) in combination



Figure 2. Geographical spread of *steadfastness*

with conditional inference trees. Both have been used for linguistic analysis in recent years (see, e.g., Tagliamonte & Baayen 2012; Hundt, Rautionaho & Strobl 2020; Röthlisberger *forthcoming* (on random forests); Ota & Green 2013; Smolíková 2014; Carlson, Sonderegger & Bane 2014; Burke, Morita-Mullaney & Singh 2016; Van de Velde & Keersmaekers 2020 (on survival analysis)). Both approaches have, to my knowledge, only rarely been applied to historical data, especially historical English data.

In the random forests method, random subsets of the entire dataset are analysed to find out which combination of variables (here: suffix, geographical spread, frequency, rhyme ratio) leads to which outcome (here: survival or obsolescence). The analysis of one such

Table 1. *Random sample of ME subset, censored at 1500*

Lexel	First att.	Last att.	Obsolete		Freq.	Rhyme freq.	Rhyme ratio	Freq. log	Spread log
			(0= survived, 1= obsolete)	Suffix					
mirkness ⁷	c.1250		0	ness	6	1	0.16667	1.79176	4.5256
rightfulness	1303		0	ness	13	1	0.07692	2.564949	5.4756
manhood	a.1225		0	hood	20	4	0.2	2.995732	5.9996
whoredom	c.1175		0	dom	36	7	0.19444	3.583519	5.9699
sickness	a.1225		0	ness	39	1	0.02564	3.663562	6.0272
wilderness	c.1200		0	ness	59	10	0.16949	4.077537	5.9029
younghood	c.1275	a.1366	1	hood	1	0	0	0	0
cwēadship	c.1205	a.1225	1	ship	6	1	0.16667	1.79176	5.4380
wretchdom	a.1225	a.1400	1	dom	6	1	0.16667	1.79176	3.8391

subset can be visualised in tree form (see e.g. [figure 17](#)), and the total number of trees generated from a large number of subsets is called a ‘forest’. Observations in the dataset that were not used in this initial stage are then used to test how accurately the forest classifies the data: the results of the individual trees are aggregated and the predicted outcome is compared to the actual, recorded outcome. The random data sample in [table 1](#) shall serve as an example to illustrate the process. First, it seems that the suffix *-ness* is a good predictor for survival because none of the words that are obsolete end in *-ness*, whereas 4 out of 6 of the surviving words do. Furthermore, in the same random sample, a high rhyme ratio seems to go hand in hand with survival, too, since the obsolete words have ratios of 16.6 per cent or lower, whereas half of the surviving words have a ratio higher than that. Similarly, the surviving words tend to have higher values for the variable Frequency: it is 1.8 or lower for the obsolete words, whereas all but one of the surviving words (*mirkness*) have a higher log frequency. Finally, only one out of three of the obsolete words in the sample (*cwēadship*) has a Spread value of equal-or-greater than 5.4, whereas the same is true for 5 out of the 6 surviving words. This leads to the preliminary assumption that words have a higher chance of surviving if they: (a) end in *-ness*; (b) have a rhyme ratio of over 16.6 per cent; (c) have a log frequency of over 1.8; and (d) have a geographical spread of over 5.4. This process is then repeated many times with different, random subsets, and the results are aggregated in order to determine the combination of criteria that are overall most effective for survival.

SAs are often used in medical studies (see examples in Kleinbaum & Klein 2012: 4–5), for example to test the effectiveness of a medication over a placebo, or to compare the survival rates of two groups (e.g. men and women). Over the course of such a study, the time up to a particular event, such as the death of a patient, is recorded. The SA is

⁷ ‘Chiefly Scottish in later use. Now rare’ (*OED*, s.v. *mirkness*). However, at the point of censoring (in this example, 1500), it would still have been in use.




Figure 3. Kaplan–Meier curves for the geographical spread in the ME subset, censoring date 1350

conducted in order to determine whether certain variables (such as type of medication, age, gender, etc.) have a significant effect. In other words, it is used to test whether certain variables influence the patient's chances of survival, i.e. how much time passes between the start of the study and the event (i.e. death). If the observation period ends before the event occurs, the data are said to be *CENSORED* at this point in time. Using [table 1](#) as an example again, *cweadship* is attested about twenty years, from 1205 to 1225. This falls within the period of observation (which in this example is, for the purpose of illustration, 1150–1500), and the event ('death' or obsolescence) occurs after twenty years. For a word like *manhood*, on the other hand, which is attested from c. 1225 to beyond the period of observation, the survival time is only measured up to 1500 (i.e. 275 years), at which point the event (obsolescence) has not yet occurred, and so it is censored.

The survival probability calculated in this manner is visualised with Kaplan–Meier (KM) survival curves such as the one in [figure 3](#), with time (in years) on the x-axis and survival probability on the y-axis. This figure shows the survival probability according to different levels of geographical spread (low, medium, high). For calculation purposes, absolute survival times are considered, that is, survival probability is calculated based on the length of attestation, regardless of when the different words first occurred during the observation period. Whenever the population or group size decreases because the event (i.e. death or obsolescence) occurred, the probability of a word in this group surviving up to this point in time decreases, and there is a step down in the 'curve'. If a word is merely censored, i.e. if the length of its attestation is

cut short not because it died out but because the time of observation ended, this is marked by a small vertical line in the curve. In [figure 3](#), there are several drops in the middle curve (representing the medium category, in green), around the 100-year mark. This indicates that there are a number of words that fell out of use approximately 100 years after they were first recorded (regardless of the specific year in which they were first recorded). In the same survival curve, there are small vertical lines that cross it, e.g. around the 50-year mark or just after the 100-year mark. This means that there are words in the medium geographical spread category which were attested during only 50 and slightly more than 100 years, respectively, during the 200-year observation period, but they do not decrease the survival probability because they did not become obsolete during this time.

With this method, only one (or very few, depending on the number of levels) categorical variable can be considered at a time. A more complex method, the Cox Proportional Hazards model, can be used to analyse several (non-categorical) variables at the same time. However, as the name implies, the model assumes that the hazard between groups is proportional, which is not the case for my dataset, as can be seen in [figures 4–15](#), where the curves often cross each other and do not maintain equal distance from each other. This, among other things, led to issues with model fit, which is why the current study only uses KM curves.

Adapting the terminology of SA to the data of the present study, words are, metaphorically speaking, the ‘patients’, the event of interest is a word’s ‘death’, for which I take the year of its last attestation, and the survival time is measured as the difference between its last and first attestation. The main variable of interest – the type of medication, if you will – is rhyme, which is measured as a word’s appearance in rhyming position in proportion to the total frequency. Additional variables that were included are the frequency, the suffix and the geographical spread. Finally, the censoring date is 1350, shortly after the timespan covered by *LAEME*. However, because of how the data are distributed, the survival curves for the French subsets were not very enlightening: only seven out of 118 words had disappeared by 1350, and most of the French loanwords in my dataset entered the language after 1225, where we find a big spike in first attestations, with another big spike around 1280. This means that there is less time for observation than in the OE and ME set (which contain words that are already attested earlier). I therefore moved the censoring date for the French subset to 1500.

Numeric variables were used for the RF, but since categorical variables were needed for the SA, I transformed them as follows: there are four frequency categories, low (1 or 2 occurrences), medium-low (3–10 occurrences), medium-high (11–31 occurrences) and high (more than 31 occurrences). The ‘low’ category is intended for words that occur only once or twice in the corpus, and the ‘high’ category for words that occur exceptionally often. A cut-off point of 31 was chosen because in both the French and the ME subset, there is a gap after this, followed by only 4 words in each subset that have a much higher absolute frequency. The geographical spread was divided into three levels: low, medium and high. The ‘low’ category includes words with a spread




Figure 4. Kaplan–Meier curves for the different rhyme ratio categories of the French subset, censoring date 1500




Figure 5. Kaplan–Meier curves for the different rhyme ratio categories of the ME subset, censoring date 1350




Figure 6. Kaplan–Meier curves for the different rhyme ratio categories of the OE subset, censoring date 1350




Figure 7. Kaplan–Meier curves for the different suffixes of the OE subset, censoring date 1350




Figure 8. Kaplan–Meier curves for the different suffixes of the ME subset, censoring date 1350




Figure 9. Kaplan–Meier curves for the different suffixes of the French subset, censoring date 1500




Figure 10. Kaplan–Meier curves for the different frequency categories of the OE subset, censoring date 1350




Figure 11. Kaplan–Meier curves for the different frequency categories of the ME subset, censoring date 1350




Figure 12. Kaplan–Meier curves for the different frequency categories of the French subset, censoring date 1500




Figure 13. Kaplan–Meier curves for the geographical spread in the ME subset, censoring date 1350




Figure 14. Kaplan–Meier curves for geographical spread in the OE subset, censoring date 1350

value of up to 4.7, which in general amounts to occurrences that are contained within one dialect area.⁸ The ‘medium’ category includes words with a spread value of up to 5.6, which is roughly equivalent to two large dialect areas. Any words with a spread value above 5.6 are in the ‘high’ category. For the rhyme ratio variable, I split the data into five categories: less than 20%, 20–40%, 40–60%, 60–80% over 80%.

The statistical tests were done in R using the *survival* (Therneau 2022) and *survminer* (Kassambara *et al.* 2021) packages for SA. For the trees and forests, I used the *party* (Hothorn *et al.* 2006; Strobl *et al.* 2007, Strobl *et al.* 2008), *partykit* (Hothorn & Zeileis 2015) and *pdp* (Greenwell 2017) packages. The RF analysis was done using conditional variable importance and the initial value of *ntree* was set to 500 and *mtry* to 2, which performed quite well for the French subset (C-index=0.923, accuracy=0.88, out-of-bag (OOB) C-index=0.79 and OOB accuracy=0.88), but not for the OE or the ME subset. Increasing *ntree* had little to no effect, but increasing *mtry* to 3 led to slightly better scores. However, the OOB C-index is still below the desired minimum of 0.7 for both ME (C-index=0.81, accuracy=0.74, OOB C-index=0.69, OOB accuracy=0.71) and for OE (C-index=0.8, accuracy=0.72, OOB C-index=0.67, OOB accuracy=0.64). There is no difference in variable importance between *mtry*=2 and *mtry*=3 for the OE subset, and only a slight increase in the importance of the Suffix variable for the ME subset with *mtry*=3. The trees were made using the default

⁸ Dialect areas: North, East Midlands, West Midlands, Southeast/Kent, Southwest.

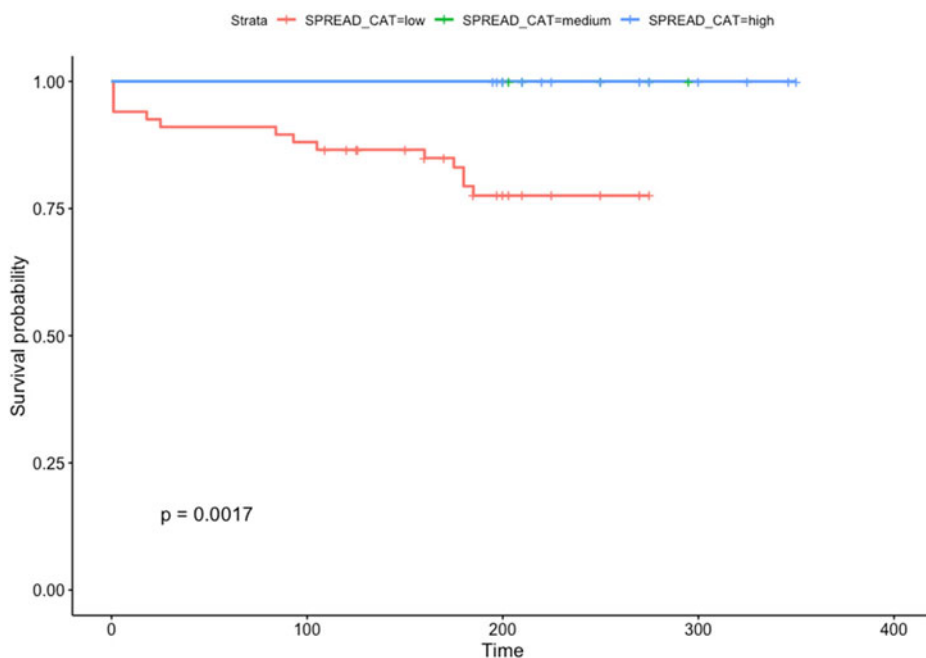


Figure 15. Kaplan–Meier curves for geographical spread in the French subset, censoring date 1500

settings (minimum p-value=.05, minsplit=20, minbucket=7). The accuracy of the trees is somewhat low, as are the C-indices: 0.69 (both accuracy and C-index) for OE, 0.68 accuracy and C-index 0.66 for ME, and 0.91 accuracy and C-index 0.67 for French. Reducing the minimum sizes for nodes (minsplit=10 and minbucket=5) somewhat improved the performance for the ME data, but not OE or French. Since the OOB C-indices for the RF and the C-indices for the trees were quite low, especially for the OE and ME subsets, the results described in the following section should be seen as exploratory.

5 Results

5.1 Survival analysis

The result of the survival analysis is visualised by the Kaplan–Meier curves in figures 4–15. The curves for the rhyme ratio variable (figures 4–6) show that there is no indication that a higher amount of rhyme leads to a higher survival probability. In all three subsets, the survival probability for the >80 per cent category falls quite steeply and/or is one of the lowest at the end of the observation period. However, rhyme nevertheless does seem to have an influence on survival, albeit only in the French and Middle English subsets – for OE, the difference between the five ratio groups is not statistically significant. In the French subset (figure 4), words with a

rhyme ratio between 20 and 60 per cent have the highest survival probability, whereas those with 60–80 per cent and more than 80 per cent rhyme ultimately have the lowest survival probability. The group of words with the lowest rhyme ratio (under 20 per cent) is somewhere in the middle all throughout the observation period. None of the five ratio groups' survival probability drops below 50 per cent during the observation period. For ME (figure 5), the picture is slightly different: here, too, the 20–40 per cent group has the highest survival probability initially, but ends up around the 75 per cent mark together with the third and fourth ratio group (40–60 per cent and 60–80 per cent). In contrast, the groups at either extreme of the rhyme–ratio spectrum fare the worst: the group with less than 20 per cent rhyme has a very low survival probability from the very beginning and both it and the group with more than 80 per cent rhyme drop below 50 per cent survival probability approximately seventy-five years after first occurrence.

Some of the curves can be explained by an imbalance in the data (see table 2). For example, in all three subsets, there is a large number of words that appear only once, which means that these automatically fall into the under-20 per cent or the over-80 per cent group, because their rhyme ratio can only be 0 or 100 per cent. Words that appear only twice also tend to fall on one or the other end of the scale. Together, words that appear only once or twice in the corpus make up 42 per cent (81/193) of the OE subset, 61 per cent (100/164) of the ME subset, and 48 per cent (57/118) of the French subset. The proportions of those that also fall into the under-20 per cent or over-80 per cent group are similar: 41 per cent (79/19) of the OE subset, 60 per cent (99/164) of the ME subset, and 43 per cent (51/118) of the French subset. In the case of the OE subset, these instances tend to be the last attestations of a word, and so the under-20 per cent and over-80 per cent survival rates are automatically a lot lower than the others. In the French subset, the under-20 per cent group contains more words with higher frequencies, so the picture looks somewhat different here. Another example is the low survival probability of the 60–80 per cent group in the French subset: this group contains only six words, half of which happen to survive. Similarly, the sudden and complete drop in the curve of the same group in the OE subset (figure 6) can be explained by its containing only one word (*lyberness*).

For the suffix variable (figures 7–9), the OE and ME subsets show slightly different behaviour. In the OE subset (figure 7), there is almost no difference between the four suffix groups in the first fifty to sixty years, after which they split in two: both *-ship* and *-ness* drop off, with *-ness* falling below the 50 per cent survival point after around 125 years (i.e. around the year 1275–1300),⁹ whereas more than 75 per cent of words ending in *-dom/-hood* are still in use by the end of the observation period (1350). In the ME subset (figure 8), all four suffix groups show a steep drop very early on, with more than 50 per cent of *-ship* words disappearing within fifteen to twenty years of

⁹ Note that, since all words in the OE set are attested at the start of the observation period, the timespan corresponds to real time (1150 + number of years on the x-axis). This is not the case for ME and French, since not all words are attested at the beginning of the observation period.

Table 2. *Absolute frequencies per language subset for the five ratio groups. Percentages given in the rightmost column are the percentages of the totals of the respective subsets.*

Subset	Rhyme category					Total
	<20%	20–39%	40–59%	60–79%	>=80%	
OE	175	10	3	1	4	193
n=1	51				3	54 (28%)
n=2	24		2		1	27 (14%)
ME	111	11	7	5	30	164
n=1	52				21	73 (44.5%)
n=2	18		1		8	27 (16.5%)
FR	52	16	20	6	24	118
n=1	19				21	40 (34%)
n=2	9		6		2	17 (14.4%)

their first attestation. The *-hood* group, which had a high survival probability in the OE subset, also falls below 50 per cent around seventy-five years after the first occurrence, while *-dom* and *-ness* have relatively similar survival curves that stay above 50 per cent until the end of the observation period. In the French subset (figure 9), almost all suffixes maintain a very high survival probability. The notable exception here are words ending in *-ery* (table 3), whose survival probability decreases quite quickly around 175–200 years after first attestation, after which it stays just above the 50 per cent mark.

The patterns for frequency (figures 10–12) and spread (figures 13–15) are quite similar, which is to be expected because they correlate (Pearson's $R = 0.8$ for OE and 0.76 for ME, $p < 0.05$ for both). The overall tendency is very clear in both subsets: words with higher frequency and geographical spread have a higher survival probability, words with low frequency and geographical spread have a lower survival probability. In the OE subset, we also see a similar pattern to that of the suffix graph: there is not much difference between the different levels of frequency/spread in the first fifty or sixty years, after which they clearly separate. There seems to be no difference between the low-medium and medium-high-frequency words. The low-frequency/spread words drop below the 50 per cent mark after about 100 years (i.e. around 1250). In the ME subset (figure 11), the survival curve for the low-medium-frequency words is very similar to that of the low frequency words, and both of them drop below 50 per cent survival probability after 100 and 75 years, respectively. In the French subset (figure 15), the words with lower geographical spread have a significantly lower survival probability than those with medium or high spread ($p = 0.0017$), but it is still relatively high compared to OE and ME. For frequency (figure 12), there are only small differences which are not statistically significant, but it is again the group with the lowest frequency (one or two occurrences) whose survival probability decreases the most.

Table 3. *Words in the dataset ending in -ery, sorted by spread. Note that those words still in use at time of censoring, i.e. 1500, have that date as their 'last attestation', but are marked as 'not obsolete' at this point.*

Lexel	First att.	Last att.	Obsolete	Total freq.	Rhyme ratio	Spread (log)
lovedruerie	1300	1405	1	1	1	0
gentlery	1275	1460	1	1	1	0
mangerie	1300	1475	1	2	0.5	0
guilerie	1303	1483	1	3	1	0
losengerie	1400	1484	1	2	0	0
bocherie	1340	1500	0	1	0	0
nunnery	1275	1500	0	4	0.75	0
druerie	1225	1500	0	10	0.5	5.633934
mastery	1225	1500	0	27	0.333333	5.710854
treachery	1225	1500	0	31	0.451613	5.823313
robbery	1200	1500	0	9	0.111111	5.887372
lechery	1230	1500	0	92	0.217391	6.07863

5.2 *Random forests and conditional inference trees*

The results of the RF analysis allow us to rank the variables according to their importance and show which variables are important and which ones are not. To complement the variable importance plots, conditional inference trees were used to visualise interactions between the variables. In some cases, partial dependence plots are used to more closely evaluate individual variables. Each of the three language subsets and their respective forests and trees are discussed in order, starting with OE and ending with French.

In the OE subset, Spread and Suffix are important predictors, while Ratio and Frequency are not (figure 16). The tree (figure 17) did not show any interactions between Spread and Suffix, and instead the importance of Spread seems to outweigh everything else. It did, however, split up the spread variable into three groups, which are the three bar-plot leaf nodes; dark grey represents the proportion of words that survived, light grey those that did not. The first split (node 1) is made between those with a spread of 5.67 and above, which have almost a 100 per cent survival rate, and those below. This group is split again (node 2) into those with a Spread value between 4.47 and 5.67, and those below 4.47. The former have a survival rate of about 50 per cent, the latter around 30 per cent. These are, coincidentally, roughly the same as the cut-off points chosen for the categorical spread variable. Suffix does not show up in the tree, even though is an important predictor according to the RF's variable importance plot – albeit much less important than Spread. A partial dependence plot for Suffix (figure 18) shows essentially the same as the KM curves: *-ness* is most likely to become obsolescent (coded as 1), followed closely by *-ship*, whereas *-hood* and *-dom* are more likely to survive (coded as 0).




Figure 16. Variable importance in the OE subset




Figure 17. Conditional inference tree for the OE subset, censoring date 1350




Figure 18. Partial dependence plot for Suffix in the OE subset. The y-axis ('yhat') represents the predicted survival probability.

For the ME subset (figure 19), too, Spread is the most important variable, followed by Suffix, while Frequency is clearly unimportant. Ratio is very close to the red line, i.e. very close to being unimportant, and a partial dependence plot (figure 20) shows that words with higher rhyme ratios tend to survive more than those with lower rhyme ratios, but the difference is very small. Ratio also does not appear in the tree (figure 21), which only includes Spread and Suffix. The first split is between those words with a Spread higher or lower than 4.362 (node 1). Those at or below this level have a roughly 30 per cent survival rate (node 2; dark grey represents surviving words, light grey obsolete words). Among the group with a higher geographical spread, the tree splits further (node 3) on




Figure 19. Variable importance in the ME subset




Figure 20. Partial dependence plot for Ratio in the ME subset. The y-axis ('yhat') represents the predicted survival probability.

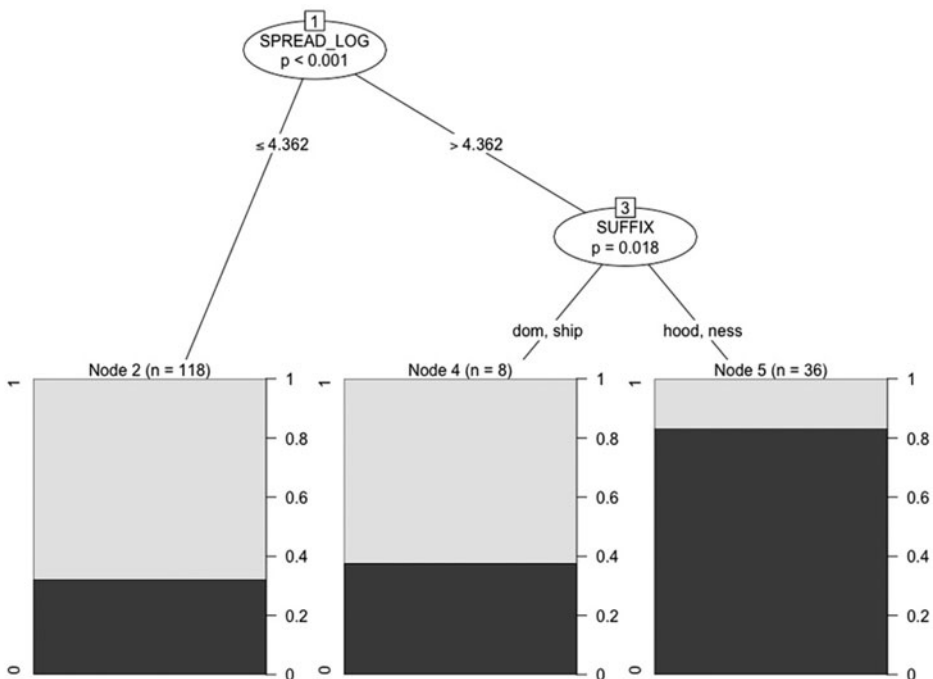


Figure 21. Conditional inference tree for the ME subset, censoring date 1350




Figure 22. Variable importance in the French subset

the Suffix variable: those ending in *-dom* or *-ship* (node 4) have a similarly low survival rate to those words with lower geographical spread, but this node contains only eight words and is therefore only applicable to a very small part of the dataset. Words ending in *-hood* or *-ness* (node 5), on the other hand, have a survival rate of just over 80 per cent.

The RF analysis of the French subset shows that both Suffix and Spread are important predictors for survival (figure 22), but unlike the OE and ME subsets, Suffix is more important than Spread. According to the tree (figure 23), those with a medium to high geographical spread (split at node 1) have a 100 per cent survival rate (node 5), but among those that fall below a certain threshold (3.748), there is another split on Suffix (node 2): those ending in *-ery* (node 3) fare decidedly worse, with a survival rate of roughly 30 per cent, than those ending in other suffixes (node 4), with over 80 per cent. This matches the survival curves for the French suffixes, which also indicated that *-ery* had a significantly lower survival probability than the other suffixes. Since the tree suggests that this may also be dependent on the geographical spread, it is worth having a look at the *-ery* group, which is given in table 3. Of the 12 words ending in *-ery*, 5 are obsolete before the end of the observation period (1500). Unfortunately, all of these are attested only between one and three times in *LAEME*, which automatically makes it very difficult, if not impossible, to gauge the extent of the geographical spread.

Overall, the results of the RF analyses and SAs complement each other. In all three subsets, geographical spread and suffix were the most important and most robust predictors for survival, usually in combination with each other. The other variables are not as important according to the RF's variable importance plot, but there nevertheless seem to be (for the most part) significant differences between the different levels of these variables, too.

6 Discussion

Finally, I would like to have a closer look at some of the cases that fall outside the expected pattern, such as words that survive despite low geographical spread, or words that fall out

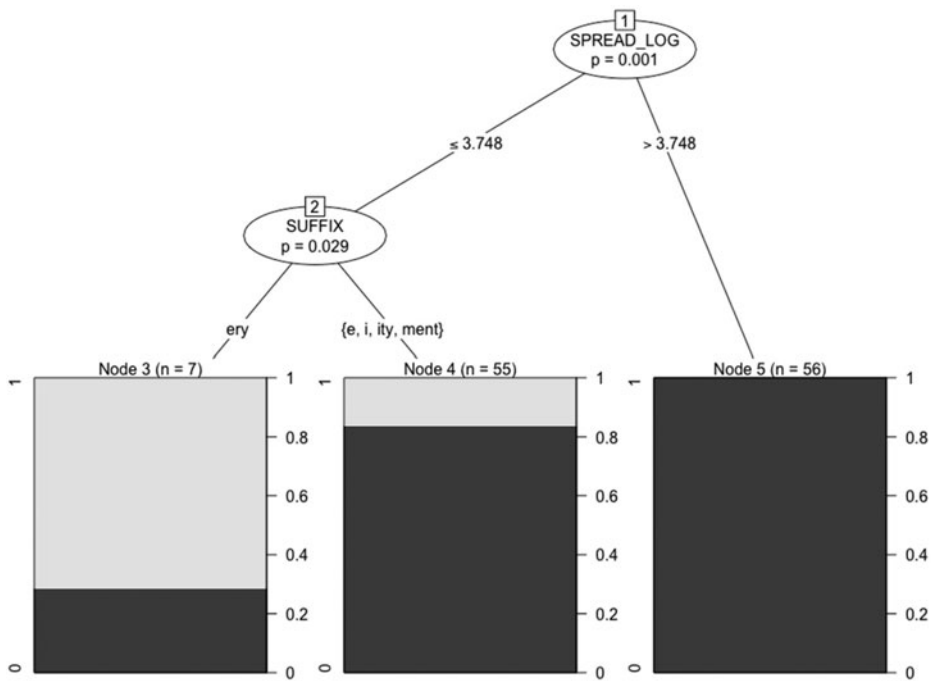


Figure 23. Conditional inference tree for the French subset, censoring date 1500

of use despite high geographical spread. The OE words that survive up to or beyond 1350 (table 4) despite an apparent disadvantage seem to be mostly deadjectival nouns, and their adjective bases are for the most part still in use today. This is the case, for example, with *soreness* and *sore*, *blindness* and *blind*, or *leanness* and *lean*. Among the obsolete words, on the other hand, we find for example *anrædnes*, *arfæstnes*, *diegelness* or *hluttornes*, which are also derived from adjectives, but their adjective bases have either already fallen out of use by the beginning of the ME period (*hlutor*, *arfæst*), or do so soon after (*anræd* c.1230, *diegol* c.1275, *HTE* s.v. *anræd*, *diegol*). A few words in this group, *bishophood*, *earldom*, *lorddom* and *lordship*, are not based on adjectives but on terms of rank or office that still exist today. One factor that is often mentioned in connection with lexical loss is that a term is more likely to fall out of use if the concept or object it denotes is no longer relevant or in use (e.g. Rundblad 2000: 40; Dworkin 1989: 335), and so the continued relevance and existence of these ranks and titles most likely contributes to their survival to this day. The trend of deadjectival nouns surviving longer alongside their adjectival bases agrees with other observations made by Rundblad (2000: 39–40), who points to the importance of transparency and the connection between the base and the derivation; that and having both the derivative and the base available in the lexicon at the same time seems to facilitate frequent use and thus survival (Rundblad 1998: 171; 2000: 40). Similarly, Dworkin (1989: 336) observes in a study on Old Spanish that ‘derivationally transparent lexical items

Table 4. *Words in the OE subset that survive up to or past 1350, where spread < 4.47. Words that have no last attestation date are still in use today.*

Level	Last att.	Total freq.	Rhyme ratio	Geographical spread (log)
earldom		2	0	0
lorddom		4	0	0
bishophood		1	1	0
blindness		1	0	0
blitheness		1	0	0
carefulness		1	0	0
dimness		2	0	0
dryness		1	0	0
fastness		1	0	0
fullness		1	0	0
greediness		2	0	0
leanness		1	0	0
mennisness		4	0	0
sharpness		2	0	0
slackness		3	0	0
soreness		1	0	0
wideness		1	0	0
wōdness		1	0	0
aceness	C13a2	2	0	0
forgetolness	c.1450	1	0	0
lithnesse		3	0	0
mēnesse	C14a2	4	0	0
swērnese	1676	1	0	0
wealdness	C14a	1	0	0
byrignese	c.1470	3	0	3.17194
dreariness ¹⁰	1596	4	0	3.831704

depend heavily on the head of the word family to which they belong ... In the case of many adjectival abstracts, the loss of the primitive (almost) inevitably led to the demise of the corresponding substantive(s).'

In the ME subset, words that have disappeared despite being attested in a wide geographical area and ending in *-hood* or *-ness* are *lordhood*, *wodhood*, *holihood* on the one hand, and *skillwiseness*, *untrueness*, *dreadness* on the other. For those ending in *-hood*, there is usually a synonym ending in *-ship*, *-dom* or *-ness* in use at the same time. In addition to *lordhood*, there are *lorddom* and *lordship* (attested since OE). In the case of *holihood*, too, *halidom* and *holiness* are attested since OE, and next to *wodhood*, *wodness* is attested since OE. While a co-existing base word and its

¹⁰ The last attestation in the sense 'state/condition of sorrow/grief' is in 1596, whereas the first attestation in the sense 'gloomy quality' is in 1727. Since these two dates are more than a hundred years apart, I assumed the newer sense to be a new, independent formation on the basis of *dreary*.

derivative may reinforce each other in terms of regularity and frequency of use (Rundblad 2000: 39–40), the same is apparently not true for the existence of parallel derivatives like *lordship/-hood/-ness* and *holihood/-dom/-ness*. Lindsay & Aronoff (2013: 135) point out that ‘in general, languages do not tolerate true synonymy’, and these synonyms seem to be synonymous enough to be sometimes interchangeable and thus in competition with each other (Gardner 2014: 29–30). Among the words ending in *-ness*, it is notable that one of the available synonyms is always the simplex from which it is derived or to which it is related: *skill* for *skillwiseness*, *untruth* for *untrueness* and *dread* for *dreadness*.

The next category in the ME subset that stands out is the group of words that survive despite a smaller geographical spread. The majority of these, too, ends in *-hood* or *-ness* and their bases generally survive(d) alongside the derivatives. In the ME period, *-ness* is ‘the most frequent suffix across the board chronologically and geographically’ (Dalton-Puffer 2011: 128). Since geographical spread was the most consistently significant predictor of survival in the SA, the general wide geographical distribution of *-ness* may thus be the reason why *-ness* is one of the factors that is associated with high survival rates in the tree (figure 21) as well as the SA (figure 8).

The one factor that stood out in the French subset is the significantly higher disadvantage of words ending in *-ery* in combination with low spread. Indeed, half of the words ending in *-ery* have become obsolete by 1500, but most of these, as mentioned above, only appear very few times in the corpus, most likely because they only appear towards the end of the period covered by *LAEME*, which often makes it impossible to determine the geographical spread. I suspect that the cause lies in the small size of the dataset and the unfortunate absence of other infrequent words ending in *-ery*. Gadde (1910: 27–31) lists several derivations in *-ery*, mostly from native English bases, from the early fourteenth century that are not in *LAEME*. Probably because his scope is larger than that of the present study or that of Gardner (2014), he also finds a much larger number of words in *-ery*, and a relatively high proportion of derivatives formed within English: ‘Out of more than 600 new-formations with *-ery* (*-ry*) in my word-list more than 250 are from native roots’ (Gadde 1910: 13). This suggests that the heyday of *-ery* only began at the beginning of the fourteenth century and that my corpus was not suited for an SA in this case.

There are few other ways in which the make-up of my dataset may have skewed the results. Due to the nature of the suffixes chosen for purely pragmatic reasons, i.e. to make data extraction easier, the words included in my dataset denote for the most part concepts that are abstract, and thus rather general (see, e.g., Trips 2009: 188, who also equates abstractness with generalness). Sylvester *et al.* (2020) looked at different levels of specificity in a handful of semantic domains and found the more general levels to have a greater number of synonyms (cf. graph 2). In other words: the more specific a sense, the more likely it is that there are only a few (one to three) co-hyponyms, or synonyms at the same level of specificity, and the more general a sense, the more likely it is that there are a lot of synonyms available. The more abstract and general the meaning of words in *-ity*, *-ness*, *-ship* etc., the more likely they are to have a large number of synonyms, and thus more competition. And while these suffixes may have

acquired rather specific meanings/semantic niches in Present-day English, they probably used to be more general at the time of borrowing (cf. Gardner 2014: 191 on *-ity* and *-ness*; Lloyd 2011: 35–9 on *-ment*).

7 Conclusion

The original aim of this article was to investigate whether frequent usage in rhyme potentially increased the survival chances of medieval English nouns. Since the focus was on rhyme, the data were chosen based on suffixes, that is, *-ship*, *-hood*, *-dom*, *-ness* for Old English words and new Middle English formations, and *-ity/-e*, *-ment*, *-ery/-i* for French loans. However, among the variables that were included, geographical spread and suffix stood out as most significant across these three subsets. For the individual subsets, there were some additional significant factors: *-ness* and *-hood* in the ME subset leading to higher survival, and *-ery* in the French subset leading to lower survival. The effect of the suffixes in the ME and French subset require more study since they may be significant only due to the limitations of the dataset and the corpus from which the data were extracted (in the case of *-ery*, which only appears towards the end of the timespan covered by the corpus) or they may be secondhand effects of another factor (overall wide geographical distribution in the case of *-ness*). The C-indices for the trees and forests in this study were also rather low, so the significance of these factors will need to be re-examined with models that perform better.

In addition to addressing these issues, future research would benefit from including additional factors which were not taken into account in the present study but stood out as potentially important in the discussion section, such as the parallel survival of the derivatives' bases and the number and specificity of available synonyms. Overall, the influence of rhyme on survival is likely to be restricted to individual words, as in the cases of *christianity* (Käsmann 1961: 99) or *onde* (Timofeeva 2020b), but does not seem to be generally measurable in the Middle English lexicon, unlike the general positive effect of geographical spread.

Author's address:

English Department
University of Zurich
Plattenstrasse 47
CH-8032 Zurich
Switzerland
johanna.vogelsanger@es.uzh.ch

References

- Burke, April M., Trish Morita-Mullaney & Malkeet Singh. 2016. Indiana emergent bilingual student time to reclassification: A survival analysis. *American Educational Research Journal* 53 (5), 1310–42.

- Carlson, Matthew T., Morgan Sonderegger & Max Bane. 2014. How children explore the phonological network in child-directed speech: A survival analysis of children's first word productions. *Journal of Memory and Language* 75, 159–80.
- Carroll, Julia M. & Margaret J. Snowling. 2001. The effects of global similarity between stimuli on children's judgment of rime and alliteration. *Applied Psycholinguistics* 22(3), 327–42.
- Coch, Donna, Giordana Grossi, Wendy Skendzel & Helen Neville. 2005. ERP nonword rhyming effects in children and adults. *Journal of Cognitive Neuroscience* 17(1), 168–82.
- Coch, Donna & Margaret M. Gullick. 2011. Event-related potentials and development. In Emily S. Kappenman & Steven J. Luck (eds.), *The Oxford handbook of event-related potential components*, 475–511. Oxford: Oxford University Press.
- Cutler, Anne, Brit Van Ooijen & Dennis Norris. 1999. Vowels, consonants and lexical activation. *Proceedings of the International Congress on Phonetic Sciences, San Francisco* 3, 2053–6.
- Dalton-Puffer, Christiane. 2011. *The French influence on Middle English morphology: A corpus-based study of derivation*. Berlin: De Gruyter Mouton.
- Dekeyser, Xavier & Luc Pauwels. 1990. The demise of the Old English heritage and lexical innovation in Middle English: Two intertwined developments. *Tijdschrift voor Germaanse Filologie* 79, 1–23.
- Durkin, Philip. 2014. *Borrowed words: A history of loanwords in English*. Oxford: Oxford University Press.
- Durkin, Philip. 2020. Contact and lexical borrowing. In Raymond Hickey (ed.), *The handbook of language contact*, 169–79. Hoboken, NJ: John Wiley.
- Dworkin, Steven N. 1989. Studies in lexical loss: The fate of Old Spanish post-adjectival abstracts in *-dad*, *-dumbre*, *-eza*, and *-ura*. *Bulletin of Hispanic Studies* 66(4), 335–42.
- Egan, Ciara, Filipe Cristino, Joshua S. Payne, Guillaume Thierry & Manon W. Jones. 2020. How alliteration enhances conceptual–attentional interactions in reading. *Cortex* 124, 111–18.
- FEW* = *Französisches etymologisches Wörterbuch*. 1948. Ed. Walther von Wartburg. Tübingen: J. C. B. Mohr.
- Gadde, Fredrik. 1910. *On the history and use of the suffixes -ERY (-RY), -AGE, and -MENT in English*. Lund: Berlingska Boktryckeriet.
- Gardner, Anne-Christine. 2014. *Derivation in Middle English: Regional and text type variation*. Helsinki: Société Néophilologique.
- Greenwell, Brandon M. 2017. pdp: An R package for constructing partial dependence plots. *The R Journal* 9(1), 421–36.
- Hothorn, Torsten, Peter Buehlmann, Sandrine Dudoit, Annette Molinaro & Mark Van Der Laan. 2006. Survival ensembles. *Biostatistics* 7(3), 355–73.
- Hothorn, Torsten & Achim Zeileis. 2015. partykit: A modular toolkit for recursive Partytioning in R. *Journal of Machine Learning Research* 16, 3905–9.
- Hundt, Marianne, Paula Rautioaho & Carolin Strobl. 2020. Progressive or simple? A corpus-based study of aspect in World Englishes. *Corpora* 15(1), 77–106.
- HTE* = *The Historical Thesaurus of English*, 2nd edn, version 5.0. 2022. Ed. Christian Kay, Marc Alexander, Fraser Dallachy, Jane Roberts, Michael Samuels & Iréné Wotherspoon. University of Glasgow. <https://ht.ac.uk/> (accessed 28 March 2022).
- Ingham, Richard. 2018. The diffusion of higher-status lexis in medieval England: The role of the clergy. *English Language and Linguistics* 22(2), 207–24.
- Käsmann, Hans. 1961. *Studien zum kirchlichen Wortschatz des Mittelenglischen 1100–1350. Ein Beitrag zum Problem der Sprachmischung* (Buchreihe der Anglia 9). Tübingen: Max Niemeyer Verlag.
- Kassambara, Alboukadel, Marcin Kosinski & Przemyslaw Biecek. 2021. *survminer: Drawing Survival Curves using 'ggplot2'*. R package version 0.4.9. <https://CRAN.R-project.org/package=survminer>

- Kleinbaum, David G. & Mitchel Klein. 2012. *Survival analysis: A self-learning text*. New York: Springer.
- Laing, Margaret. 2007. Guide to the index of sources. *A Linguistic Atlas of Early Middle English, 1150–1325*, version 3.2. Edinburgh: The University of Edinburgh. www.lel.ed.ac.uk/ihd/laeme2/guideToIndex.html (accessed 20 December 2022).
- Laing, Margaret. 2013–. *A Linguistic Atlas of Early Middle English, 1150–1325*, version 3.2. Edinburgh: The University of Edinburgh. www.lel.ed.ac.uk/ihd/laeme2/laeme2.html (accessed 20 December 2022).
- Laing, Margaret & Roger Lass. 2007. Chapter 4: Tagging. *A Linguistic Atlas of Early Middle English, 1150–1325*, version 3.2. Edinburgh: The University of Edinburgh. www.lel.ed.ac.uk/ihd/laeme2/laeme_intro_ch4.html (accessed 28 March 2022).
- Lea, R. Brooke, David N. Rapp, Andrew Elfenbein, Aaron D. Mitchel & Russell Swinburne Romine. 2008. Sweet silent thought: Alliteration and resonance in poetry comprehension. *Psychological Science* 19(7), 709–16.
- Lindsay, Mark & Mark Aronoff. 2013. Natural selection in self-organizing morphological systems. In Nabil Hathout, Fabio Montermini & Jesse Tseng (eds.), *Morphology in Toulouse: Selected proceedings of Décembrettes 7*, 133–53. Munich: Lincom.
- Lloyd, Cynthia. 2011. *Semantics and word formation: The semantic development of five French suffixes in Middle English*. Oxford: Peter Lang.
- Mitra, Priya & Donna Coch. 2018. An ERP study of cross-modal rhyming: Influences of phonology and orthography. *Psychophysiology* 56(4), 1–16.
- OED = *Oxford English Dictionary online*. 2016. Oxford University Press. www.oed.com (accessed 20 December 2022).
- Ota, Mitsuhiro & Sam J. Green. 2013. Input frequency and lexical variability in phonological development: A survival analysis of word-initial cluster production. *Journal of Child Language* 40(3), 539–66.
- Pascual, Rafael J. 2017. Oral tradition and the history of English alliterative verse. *Studia Neophilologica* 89(2), 250–60.
- Rapp, David N. & Arthur G. Samuel. 2002. A reason to rhyme: Phonological and semantic influences on lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28(3), 564–71.
- Röthlisberger, Melanie. Forthcoming. Macro- and micro-level variation in the English dative alternation: The view from World Englishes. In Eva Zehentner, Melanie Röthlisberger & Timothy Coleman (eds.), *Ditransitive constructions in Germanic languages: Diachronic and synchronic aspects*. Amsterdam: John Benjamins.
- Rundblad, Gabriella. 1998. *Shallow brooks and rivers wide: A study of lexical and semantic change in English nouns denoting 'watercourse'* (Acta Universitatis Stockholmiensis 88). Stockholm: Almqvist & Wiksell International.
- Rundblad, Gabriella. 2000. On the correlation between lexical stability and word creation device. *Journal of Quantitative Linguistics* 7(1), 31–41.
- Scragg, Donald G. 2013. The nature of Old English verse. In Malcolm Godden and Michael Lapidge (eds.), *The Cambridge companion to Old English literature*, 50–65. Cambridge: Cambridge University Press.
- Skaffari, Janne. 2017. Chapter 10: Language contact: French. In Laurel J. Brinton & Alexander Bergs (eds.), *The history of English: Middle English*, vol. 3, 184–204. Berlin and Boston: De Gruyter Mouton.
- Slowiaczek, Louisa M., James M. McQueen, Emily G. Soltano & Michelle Lynch. 2000. Phonological representations in prelexical speech processing: Evidence from form-based priming. *Journal of Memory and Language* 43(3), 530–60.

- Smolík, Filip. 2014. Noun imageability facilitates the acquisition of plurals: Survival analysis of plural emergence in children. *Journal of Psycholinguistic Research* 43(4), 335–50.
- Strobl, Caroline, Anne-Laure Boulesteix, Achim Zeileis & Torsten Hothorn. 2007. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics* 8(25).
- Strobl, Caroline, Anne-Laure Boulesteix, Thomas Kneib, Thomas Augustin & Achim Zeileis. 2008. Conditional variable importance for random forests. *BMC Bioinformatics* 9(307).
- Sylvester, Louise, Megan Tiddeman & Richard Ingham. 2020. An analysis of French borrowings at the hypernymic and hyponymic levels of Middle English. *Lexis* (16). <http://journals.openedition.org/lexis/4841> (accessed 2 March 2022).
- Sylvester, Louise, Megan Tiddeman & Richard Ingham. 2021. Lexical borrowing in the Middle English period: A multi-domain analysis of semantic outcomes. *English Language and Linguistics* 26(2), 237–61.
- Tagliamonte, Sali A. & R. Harald Baayen. 2012. Models, forests, and trees of York English: *was/were* variation as a case study for statistical practice. *Language Variation and Change* 24(2), 135–78.
- Therneau, Terry. 2022. *A package for survival analysis in R*. R package version 3.4-0. <https://CRAN.R-project.org/package=survival>
- Timofeeva, Olga. 2020a. On the margins of Bible translation: English decalogues and their circulation in the thirteenth–fourteenth centuries. *Aevum* 94(2), 317–40.
- Timofeeva, Olga. 2020b. *Nithe and onde, þat is here broþer*: Lexicalisation of a Middle English binomial. In Rodrigo Pérez Lorigo, Carlos Prado-Alonso & Paula Rodríguez-Puente (eds.), *‘Of ye olde English langage and textes’: New perspectives on Old and Middle English language and literature*, 51–80. Berlin: Peter Lang.
- Trips, Carola. 2009. *Lexical semantics and diachronic morphology: The development of -HOOD, -DOM and -SHIP in the history of English*. Tübingen: Max Niemeyer Verlag.
- Van de Velde, Freek & Alek Keersmaekers. 2020. What are the determinants of survival curves of words? An evolutionary linguistics approach. *Evolutionary Linguistic Theory* 2(2), 127–37.