Analogy in Suffix Rivalry - the Case of English *-ity* and *-ness*

SABINE ARNDT-LAPPE

Heinrich-Heine-Universität Düsseldorf

in print, to appear in:

English Language and Linguistcs, November 2014

Abstract

Rivalry between the two English nominalising suffixes *-ity* and *-ness* has long been an issue in the literature on English word-formation (cf. esp. Marchand 1969, Aronoff 1976, Anshen & Aronoff 1981, Romaine 1983, Riddle 1985, Giegerich 1999, Plag 2003, Säily 2011, Baeskow 2012, Lindsay 2012, Bauer et al. 2013: chpt. 12). Both regularly attach to adjectival bases, producing nouns with (mostly) synonymous meanings. Most standard accounts assume that stronger restrictiveness of *-ity* is an effect of *-ity* being less productive than *-ness*, and that the observed preferences are an effect of selectional restrictions imposed on bases and/or suffixes. The focus of the present study is on the productivity of the two suffixes in synchronic English and on the diachronic development of that productivity in the recent history of the language. The paper presents a statistical analysis and a computational simulation with an analogical model (using the AM algorithm, Skousen & Stanford, 2007) of the distribution of *-ity* and *-ness* in a corpus comprising some 2,700 neologisms from the *Oxford English Dictionary* from three different centuries (the 18th, 19th, and 20th century). Statistical analysis of the OED data reveals that *-ity* preference for pertinent bases is far more widespread than hitherto thought. Furthermore, the earlier data show a consistent development of these preference patterns over time. Computational modelling shows that AM is highly successful in predicting the variation in synchronic English as well as in the diachronic data solely on the basis of the the formal properties of the bases of nominalisation. On the basis of a detailed analysis of the AM model it is shown that, unlike many previous approaches, an analogical theory of word-formation provides a convincing account of the observed differences between the productivity profiles of the two nominalising suffixes and their emergence over time.

1 INTRODUCTION[1]

Examples of *-ity* and *-ness* derivatives that the *Oxford English Dictionary* (OED) records as 20th century neologisms are provided in (1). [2]

(1) (a) connectability (♦ connectable)
nordicity (♦ nordic)
metaphoricity (♦ metaphoric)
prescriptivity (♦ prescriptive)
picayunity (♦ picayune)

(b) excitingness (♦ exciting)
genericness (♦ generic)
blokishness (♦ blokish)
commutativeness (♦ commutative)
norseness (♦ norse)

Two empirical facts have posed a challenge to researchers. The first concerns the contemporary distribution of the two nominalisers in Modern English, which is neither complementary nor fully random. For example, both *-ity* and *-ness* are found to attach to morphologically simplex adjectival bases and to bases ending in Latinate suffixes (e.g. *connectable* ♦ *connectabil-ity*, *accept-able* ♦ *acceptable-ness*). By contrast, only *-ness* combines with adjectival bases ending in Germanic suffixes (e.g. *exciting* ♦ *exciting-ness*). It thus seems that *-ness* is less restrictive than *-ity* in the sense that it can attach to more different bases. In most descriptions of English word-formation this is interpreted as a difference in productivity between the two suffixes. Intriguingly,

[1] Many thanks are due to audiences at conferences for useful discussion of earlier versions: *Jahrestagung der Deutschen Gesellschaft für Sprachwissenschaft*, Potsdam, March 2013. *New Territories in Word-Formation*, Sofia, May 2013. Workshop *Analogy - Copy - Representation*, Bielefeld, November 2013. Masterclass *Analogy in English Word-Formation*, University of Manchester, March 2014. I owe special thanks to Ingo Plag, Melanie Bell, and Gero Kunter for lots of constructive feedback, as well as to two anonymous ELL referees and the editor Bernd Kortmann for their useful and encouraging comments. Finally, thanks to several people who have provided technical help in various stages: Linda Hilkenbach, Julia Homann, Lea Kawaletz, Alexander Kerber, Vineeta Michallow, and the students of my 2011 *Analogy* seminar at the University of Siegen. Needless to say, all remaining errors are my own.

[2] Note that *-ity* and *-ness* are not the only suffixes that can be used to nominalise adjectival bases in English. Other suffixes are *-th* (*depth* ♦ *deep*), *-cy* (*bankruptcy* ♦ *bankrupt*), and, potentially, *-ism* (*capitalism* ♦ *capitalist*). Nothing will be said about the status and role of these suffixes here. Cf. Bauer et al. (2013): chpt. 12 for some discussion.

however, it has also often been noted that for some bases to which both *-ity* and *-ness* can attach in principle, *-ity* constitutes the preferred option (e.g. Anshen & Aronoff 1981, Romaine 1983, Lindsay 2012), and it is as yet unclear how this is compatible with the above observation. Nor is it clear empirically which types of bases indeed show a preference towards *-ity*. An uncontroversial case is the behaviour of adjectival bases ending in *-able* in contemporary English, for which it has been shown in both experimental and corpus-based studies that speakers clearly prefer *-ity* as a nominalising suffix (Anshen & Aronoff 1981, Romaine 1983, Lindsay 2012). For other bases the situation is less clear.

The second empirical fact that has attracted attention in the literature concerns the diachronic development of the two nominalising suffixes. Etymologically, the pair constitute a showpiece example of the stratification of English derivational morphology. *-ness* is a Germanic suffix, which has been productive in English since Old English times. *-ity*, by contrast, is a Romance suffix, which has entered the English language during the Middle English period (about the 14th century, cf. Riddle 1985: 451). Both suffixes have been used productively in English since then (cf. esp. Riddle 1985). The central question that emerges, then, is how productivity of *-ity* evolved, and how the development of *-ity*'s productivity diachronically affected that of its rival *-ness*. Here recent research has worked under the assumption that one driving force in the development of rival affixes is synonymy avoidance. On the basis that *-ity* and *-ness* can indeed be synonymous in English (cf. Riddle 1985, cf. Bauer et al. 2013: chpt. 12, Baeskow 2012 for evidence for tendencies of semantic differentiation), Mark Lindsay (Lindsay 2012, Lindsay & Aronoff to appear) has recently proposed an interesting model which incorporates formal, morphological properties of potential bases into a model of the development of the rivalry (cf. section 2.2 below for detailed discussion). The model's historical focus is on Middle English, in which period Lindsay considers the massive influx of loanwords ending in *-ity* to be the major trigger of the development. For later periods in history, however, it is empirically unclear how the rivalry developed. Furthermore, the model proposed by Lindsay is vague about what the predictions would be for such later stages.

Based on the two empirical challenges outlined above, rivalry between *-ity* and *-ness* provides an interesting testing ground for current theories of productive

word-formation. Whereas some theories in the generative tradition have worked from the assumption that productivity of a morphological process is inversely correlated with lexical listedness (for pertinent claims about *-ity* and *-ness* cf. e.g. Aronoff 1976, Embick & Marantz 2008), many recent approaches have incorporated a lexical mechanism that generates word-formation patterns that are to some extent productive (often termed an associative component, cf. e.g. Anshen & Aronoff 1988, Anshen & Aronoff 1999, Bermúdez-Otero & McMahon 2006, Bermúdez-Otero 2012 for pertinent claims about *-ity*). In both strands of theories, *-ity* is expected to be less productive than *-ness*. The fact, however, that *-ity* is preferred to *-ness* in certain domains does not emerge in a straightforward way. Nor do these theories say much about diachronic change in the relative distribution of the two rival suffixes.

One class of theories that seems to be particularly suited to doing just that (i.e. predicting more finegrained productivity differences and predicting diachronic change) is found in analogy-based theories of word-formation. Such theories assume that affix selection happens online based on pertinent distributions in the mental lexicon. Variants are exemplar-based or analogical theories as proposed, for example, in Derwing and Skousen (1989), Skousen (1989 et seq.), or Daelemans et al. (2005 et seq.). Crucially, such theories do not embrace a stratal distinction between a rule-based and an associative component of grammars. Furthermore, they are based on the assumption that derived words are available as stored instances in the lexicon. This type of approach has not been tested on *-ity* and *-ness* rivalry. However, there is quite an extensive body of research that shows that such models are capable of predicting different productivity profiles of rivals or allomorphs. Much of the literature has concentrated on English past-tense formation (cf. esp. Eddington 2000, Keuleers 2008, Chandler 2010 for recent contributions), but there is also a growing body of literature concerned with phenomena in word-formation (cf. e.g. Chapman & Skousen 2005, Eddington 2006, Arndt-Lappe & Bell 2013). Furthermore, as Chapman & Skousen (2005) have convincingly shown for negative prefixation in English, an analogical approach is also capable of modelling diachronic dynamics in affix rivalry. However, the exact details about what exactly is the basis for the prediction of change still remain somewhat opaque.

The model proposed by Lindsay (2012) and Lindsay & Aronoff (to appear) is difficult to place among the groups of models outlined above. The reason is that

pertinent texts are vague about what mechanism they exactly assume to be operative. Thus, Lindsay suggests that the synchronic productivity of *-ity* in specific morphological classes is an effect of principles that can be likened to those of natural selection in biology. Due to a perturbation of the system that was caused by the massive influx of Romance words ending in *-ity* in late Middle English, *-ity* was able to become productive in niches. Apart from the metaphor of natural selection, however, it is unclear what mechanism decides on which affix is selected and, in particular, which mechanism is responsible for the retreat of *-ness* in certain domains. A potential role of exemplars is indeed mentioned in the text (e.g. Lindsay 2012: 192, 196, Lindsay & Aronoff to appear: sections 1.1, 1.2,), but the theory is not spelled out.

The present paper has two aims, one empirical and one theoretical. Empirically, the paper seeks to contribute to a better understanding of both the contemporary and the diachronic division of labour between *-ity* and *-ness*. I will present a quantitative analysis of the distribution of *-ity* and *-ness* among neologisms found in the *Oxford English Dictionary*, focussing on issues concerning both the productivity and the historical development of the two nominalisers. For productivity, the study will look at all pertinent 20th century neologisms. Results largely confirm the findings of Lindsay's recent corpus-based analysis, but suggest that *-ity* preference with certain bases seems even more widespread than documented in Lindsay's study. To assess diachronic development, the study will look at neologisms in the two centuries leading up to present day (the 18th, 19th, and 20th centuries). The analysis will thus complement earlier research, which has mainly focussed on Middle English as a crucial period. It will be shown that the development that has begun with the introduction of *-ity* in Middle English exhibits an ongoing dynamics also in later periods. Specifically, we observe a consistent diachronic extension of *-ity* usage into domains that have hitherto been dominated by *-ness*.

The theoretical aim of the paper is to explore in how far *-ity* / *-ness* rivalry can be accounted for in an analogical model. AM (Skousen & Stanford 2007) will be used as a computational implementation to predict both the synchronic distribution of the two nominalisers and their diachronic development. We will see that, indeed, AM is able to successfully predict the distribution in all experiments, soleley on the basis of the phonologial makeup of the two base-final syllables and some rudimentary syntactic

information. A detailed analysis of the model will be provided to see which predictions exactly are made concerning pertinent lexical distributions and mechanisms. With respect to productivity, we will see that the different productivity profiles of the two nominalisers emerge from differences in terms of the bases of pertinent analogies. *-ity* is assigned on the basis of more local analogies than *-ness*. This is fully in line with findings from work in analogical modelling of other morphological phenomena (discussed in detail esp. in Keuleers 2008, Arndt-Lappe & Bell 2013). With respect to diachronic development, the model predicts that two aspects are especially relevant for change. These are the (relative) locality of the generalisation and the probabilistic nature of the prediction. Whereas overall the model makes good predictions for the diachronic data, it remains a task for future research to test whether these predictions are borne out.

The paper is structured as follows: Section 2 will review current approaches to *-ity* and *-ness* rivalry, focussing on problems of productivity and diachronic development. Section 3 will then provide a quantitative analysis of the OED data in light of the issues raised in section 2. Sections 4 and 5 will be concerned with the theoretical proposal: The analogical model will be introduced in section 4, section 5 will present the results of the simulation studies with the OED data. The paper ends with a concluding section (section 6).

2 CURRENT APPROACHES TO *-ITY* / *-NESS* RIVALRY

2.1 *Synchronic accounts: degrees of productivity*

The productivity of an affix is taken to mean here, in a very general sense, the likelihood that new words are coined with that affix. Scholars generally agree that *-ity* and *-ness* are both productive in contemporary English. At the same time, however, it is often pointed out that there is a productivity difference between them, with *-ness* being the more productive nominaliser overall. Several issues, some of which are theory-dependent, bear on the question how productivity of *-ity* and *-ness* is assessed in the literature.

Thus, in stratum-based generative theories, lower productivity of *-ity* is expected given its stratal affiliation (esp. Kiparsky 1982a, Kiparsky 1982b, Giegerich 1999). *-ity* is a class 1 suffix, *-ness* is a class 2 suffix. This is evidenced not only by differences in

productivity, but also by a variety of morphophonological alternations that are found with *-ity*, but not with *-ness*. A common segmental alternation is velar softening (Chomsky & Halle 1968), where base-final [k] alternates with [s] in the derivative as in the pair *electri*[k] ⬩ *elektri*[s]*ity* and *opa*[k] ⬩ *opa*[s]*ity*). On the prosodic level, *-ity* suffixation causes stress shift to the antepenultimate syllable of the derivative (cf. Liberman & Prince 1977, Hayes 1982). Examples are pairs like *nordícity* and *nórdic* or *prescriptívity* and *prescríptive*. Another metrical effect is trisyllabic laxing (Chomsky & Halle 1968), an optional process where base-final tense vowels are changed into lax vowels (as in *ser*[i:]*ne* ⬩ *ser*[e]*nity* vs. *ob*[i:]*se* ⬩ *ob*[i:]*sity*; cf. e.g. Bauer et al. 2013: chpt. 9 for discussion of the optionality of the process).

A second issue that bears on productivity is that, in terms of the bases to which the two suffixes can attach, the distribution of *-ity* and *-ness* is semi-systematic. Discussions of this fact in the theoretical literature have almost exclusively focussed on words derived from suffixed adjectival bases, for which it is commonly assumed that the morphological category of the base is relevant for the selection of the nominalising suffix. A systematic difference between *-ity* and *-ness* in this respect is that bases with Germanic suffixes exclusively take *-ness*, whereas bases with Latinate suffixes show variability between *-ity* or *-ness*. The context in which most pertinent discussion is found is the literature on affix ordering, modelled in terms of selectional restrictions (esp. Fabb 1988, Plag 1996, Plag 1999) or lexical strata (esp. Kiparsky 1982a, Kiparsky 1982b, Giegerich 1999). Apart from the categorial, hence systematic distinction between Germanic and Latinate bases, however, the distribution of *-ity* and *-ness* is semi-systematic in the sense that, within the group of adjectival bases which allow *-ity*, there are preference patterns that seem to depend on the morphological category of the adjectival base. A seminal study to show this is Anshen and Aronoff's (1981) experimental investigation of nominalisations based on adjectives ending in *-able*/*-ible*, where subjects in an acceptability judgement task consistently preferred *-ity* over *-ness* for such bases. For other adjectival classes, especially those ending in *-ive* and *-ous*, preferences are debated in the literature (cf. Romaine 1983, Anshen & Aronoff 1988, Baayen & Renouf 1996, Anshen & Aronoff 1999). In a recent corpus-based study of rival pairs with *-ity* and *-ness* that is based on Estimated Total Hits (ETM) values in Google, Lindsay (2012) finds strong support for the non-uniform productivity of *-ity*

and *-ness*. In particular, there is a sharp contrast between, on the one hand, a clear preference for *-ness* (reflected in higher ETM frequencies for pertinent *-ness* words) on a general level, and, on the other hand, a clear preference for *-ity* with certain types of bases. These types of bases have commonly been described in terms of morphological classes. These are adjectives ending in *-able*, *-al*, *-ic*, and *-ar*. Lindsay refers to these bases as 'niches', in which *-ity* is productive.

Another factor that has been argued to influence productivity of *-ity* and *-ness* is the syntactic category of the base. Again the distribution is semi-systematic. Whereas both suffixes predominantly attach to adjectival bases, the literature also reports cases where the suffixes attach to other bases. This is particularly true for *-ness*, which may occur on nouns (e.g. *appleness*), phrases (e.g. *at-homeness*), or other categories (*aboveness*, *alwaysness*, *there-ness*; all examples are from Bauer et al. 2013: chpt. 12). *-ity*, by contrast, seems to be more selective in terms of the word-class of the base. However, also *-ity* can occur with non-adjectival bases. A frequent type are bound bases (e.g. *disparity*, *amenity*), but the literature also documents cases of nominal bases (e.g. *nerdity*, *scholarity*, *spherity*, all examples are from Bauer et al. 2013: chpt. 12).

Finally, studies employing quantitative productivity measures have shown that, also in terms of established corpus-based productivity measures, *-ness* seems to be more productive than *-ity*. In particular, *-ness* is usually found to have a larger number of hapaxes in corpora than *-ity* (cf. esp. Baayen & Renouf 1996, Plag 2006). Note that this holds for an overall comparison of the two nominalisers, not for the cases of localised productivity outlined above.

In sum, the productivity profiles of the two rival suffixes *-ity* and *-ness* are interesting for an account of suffix competition. Thus, we see that *-ity* is productive, but in general more restricted than *-ness*. Moreover, however, *-ity* suffixation is subject to restrictions of different degrees. There are domains where it seems to be excluded categorically (such as adjectival bases ending in Germanic suffixes), domains where it is highly productive (such as adjectival bases ending in *-able* and bound bases), and domains where its productivity seems to be subject to negotiation with its competitor *-ness*. *-ness*, by contrast, is globally productive in the sense that, with the notable exception of bound bases, it is never excluded categorically, but can attach, in

principle, to almost any eligible base. In their discussion of nominalising suffixes in English, Bauer et al. (2013) thus come to the conclusion that

> *-ness* seems in effect to serve as a sort of default way of forming abstract nouns from non-verbal categories in contemporary English.
>
> (Bauer et al. 2013: 246)

In other domains where *-ity* is productive, however, productivity of *-ness* is restricted.

The literature on the productivity of *-ity* and *-ness* also raises questions that are important for a general understanding of the competition. These especially concern the empirical foundations. First of all, the discussion in the literature has largely concentrated on specific domains, whereas for other domains the situation is unclear. Some domains, such as unsuffixed adjectival bases and nominal bases, are hardly ever addressed in discussions of the rivalry. For other domains, preference patterns are debated. Whereas, for example, it is pretty uncontroversial that *-ity* is the predominant nominaliser for bases ending in *-able*, for other bases preference patterns are largely unclear.

A further problem is that the few empirical studies that are available are based on existing words. This holds for Lindsay's (2012) corpus study as well as Romaine's (1983) experimental study. An obvious problem for this type of data is that, given that it is well known that outputs of derivation may be stored in the lexicon regardless of the degree of productivity of the suffix, it is difficult to control for the effect of established, lexicalised words.

Other questions concern the status of pertinent domains as morphological domains. In other words, the question is whether *-ity* or *-ness* selection is based on the morphological category of its base or whether it is based on purely formal aspects of the base. The former assumption is the one that is tacitly assumed to be true in most of the literature. However, at the same time there exists a large body of literature that shows that derivational affixation in general is subject to purely formal constraints or patterns. On the one hand, many selectional restrictions in affixation have been shown to be best captured in terms of output-oriented phonological constraints (cf. esp. Plag 1999, Raffelsiefen 1999, Zamma 2012 for English). This crucially also holds for effects which have traditionally been descibed with the help of morphological domains (cf. esp. Plag

1999: 190ff. on stem truncation). On the other hand, it has been shown also outside the constraint-based literature that formal characteristics of bases often fully suffice to explain affix selection (cf. e.g. Chapman & Skousen 2005, Eddington 2006, Lindsay & Aronoff to appear on *-ic* vs. *-ical*[3] and *-ize* vs. *-ify*). Often no reference to the meaning-based aspect of morphological categories is needed. For *-ity* and *-ness*, it seems at present impossible to say whether pertinent domains are best described in terms of morphological categories or in terms of only the form of the base. The reason is that evidence that would enable us to distinguish between these two types of domain is currently lacking. Such evidence would, for example, show that either of the two suffixes is or is not selected if the base formally resembles a morphological category but does not belong to that category. A potential piece of evidence in favour of a form-based view could be found in words like *perspectivity* (♦ *perspective*, n., OED, 20th century). As to be shown in this paper, *-ity* constitutes the preferred rival for bases ending in *-ive* among 20th century neologisms. In the base *perspective*, however, the string [ɪv] does not have a morphological status. Given the many unresolved issues about preference patterns for *-ity* and *-ness* (cf. above), however, it would be premature to draw conclusions from this one datum.

A final issue that has not been resolved yet is the question of the relevance of meaning-related factors. These can be semantic (cf. esp. Riddle 1985, Plag 2003: 66–67, Baeskow 2012) or related to pragmatic and register-related factors (cf. esp. Guz 2009, Säily 2011, Bauer et al. 2013: 245ff.). With the exception of Säily's and Guz's corpus-based studies of gender and register influences, evidence for semantic or pragmatic differences between the two nominalisers has been based so far on isolated examples, systematic empirical studies are lacking. What has emerged from existing work is that differences are likely to exist in terms of pragmatic factors rather than in terms of denotation (cf. Bauer 2013r: 257ff. for discussion).

[3] Strictly speaking, Lindsay & Aronoff (to appear) use the term 'morphological niche' to describe the effects found for *-ic* and *-ical*. However, in their analysis they find neighbourhood effects purely based on formal similarity.

2.2 *The diachronic dimension*

Research on the diachronic development of rivalry between *-ity* and *-ness* has so far largely concentrated on the effects of the introduction of *-ity* in the Middle English period and its successive establishment as a productive suffix. In particular, there is an ongoing debate about the role of loanwords, borrowings, calques, and their morphological integration (cf. Dietz in press for a comprehensive recent summary that is not specific to *-ity*).

A recent approach that tries to account for the development is the one proposed by Lindsay and Aronoff (Lindsay 2012, Lindsay & Aronoff to appear). On the basis of a series of recent empirical studies of derivational suffixes in English based on data from the OED, they argue that natural selection processes are an adequate metaphor to describe the development of Romance suffixes in English: After their introduction as parts of loanwords, such suffixes have occupied 'niches', which are defined in terms of potential bases, and their fate as productive or unproductive suffixes in the course of history depends on the availability of pertinent bases in the lexicon. For *-ity*, Lindsay proposes that the increase in productivity of *-ity* was triggered by the influx of Romance loanwords in late Middle English, and his analysis seems to imply that the types of bases for which *-ity* suffixation became productive subsequently, were those where such loanwords were particularly large in number so that speakers were able to generalise. One problem that this approach faces is to explain how *-ness* lost ground as compared to *-ity* in the niches in which *-ity* became productive. The explanation suggested by the authors is the avoidance of synonymy, which they consider to be the "primary driving force behind competition in the lexicon of a language" (Lindsay & Aronoff to appear: 2). The idea that type blocking is a pervasive mechanism in the English derivational system, however, is controversial (cf. e.g. Plag 1999: 227ff., Bauer et al. 2013: 575ff. for detailed discussion of counterevidence).

A different approach is taken by Riddle (1985). Contrary to Lindsay and Aronoff, she proposes that integration of *-ity* into the English morphological system involved semantic differentiation. In her view, the process was triggered by the high proportion of *-ity* words denoting (often abstract) entities that came with French and Latin loanwords in specific areas of language use (most notably religion). As a consequence, she argues, *-ness* lost its entity sense in many cases, for which *-ity* came to

be the dominant marker. The development of the rivalry following Middle English is then interpreted as involving a process of lexical diffusion. Crucially, however, she assumes that the semantic difference is subtle, and that other, mostly sociolinguistic factors played a role as well. It is therefore not entirely clear how her view is compatible with the evidence cited in other publications that synchronically, *-ity* and *-ness* derivative pairs are often synonymous (cf. esp. Bauer et al. 2013: chpt. 12). Furthermore, it is unclear how the semantic phenomena reported by Riddle are related to the niche behaviour of the productivity of *-ity* that is reported by Lindsay (2012).

Apart from the cited approaches that conceptualise the effects of *-ity*'s integration into the morphological system of English, not much is known about the development of the division of labour between the two suffixes in the history of the language. This is particularly true for later periods, for which we find general remarks in the literature about the role of wellknown sociological factors in lexical change, but where systematic empirical studies of *-ity* and *-ness* formation are scarce (an exception is Säily & Suomela's 2009 study on gender influences in Early Modern English). An interesting period in this respect is the recent history of the language, specifically the period known as Late Modern English (18th - early 20th century; recent surveys can be found e.g. in Beal 2004, Beal 2012, van Tieken-Boon Ostade 2009). The chief reason is that Late Modern English is uncontroversially a period of large-scale lexical expansion and change (for an overview cf. e.g. Dossena 2012). It has often been claimed on general grounds that this has had consequences also for the development of the productivity of word-formation patterns, which especially affected Latinate affixes in the learned vocabulary. Systematic empirical studies which test this assumption, however, are lacking.

3 Synchronic productivity and its recent history: The OED data

In this section the issues discussed in the present paper will be investigated on the basis of neologisms gathered from the OED for three centuries (the 18th, the 19th, and the 20th centuries). The focus of the analysis will be on the 20th century situation, and the

diachronic data will be used to analyse the emergence of that situation in the two earlier centuries.[4]

The purpose of the analysis is to assess the empirical basis of the claims discussed in section 2 about differences between the productivity profiles of the two rival suffixes. In particular, the study will investigate in how far major factors that have been used in the literature to describe differences between domains for *-ity* and *-ness* affixation are reflected in the pertinent relative distributions of *-ity* and *-ness* derivatives in the OED. These factors are (a) the syntactic category of the base, (b) the morphological category of adjectival bases, and (c) the semantics. The analysis will be monovariate, comparing the proportions of *-ity* and *-ness* derivatives in the relevant subsets. The reason why only a monovariate analysis is conducted is that, as we will see, the morphological category of adjectival bases is the only parameter that produces a convincing split in the data. Syntactic category information, by contrast, is in these data only relevant in terms of a very basic distinction between bases that have word status and those that do not and are either bound forms or larger constructs. Semantic information, which can be coded here only in terms of a very rudimentary classification, does not seem to play a role. The analysis will pave the way for the analogical model to be described in section 4, which uses only formal aspects of the base and a basic word class specification as its information source.

The section is structured as follows. The extraction of data and the coding will be described in section 3.1. Section 3.2 will be concerned with productivity in the 20th century data, section 3.3 will deal with the Late Modern English data.

3.1 *Data and coding*

The database comprises the full set of *-ity* and *-ness* derivatives for which the OED records a date of first attestation between the 18th and the 20th century. Data collection involved extraction of all headwords ending in the strings 'ity' and 'ness' (using the 'Advanced Search Mode' of the OED online edition[5]) and the manual cleaning of pertinent lists. In the course of the latter procedure, an effort was made to remove all

[4] Periodisation into portions defined by centuries is - admittedly - very broad, for purely practical reasons. It is, however, sufficient to outline the general trends of development.
[5] Last accessed: Feb. 2013.

items for which it was not clear that they were transparent *-ity* / *-ness* derivatives at the time of their first attestation. In particular, this meant the exclusion of words which have entered the language as whole-word borrowings. This was assumed to be the case if (a) the OED lists a foreign word as the word's etymon, and if (b) the base of *-ity* / *-ness* suffixation is not attested in the OED earlier than the derivative. The total number of derivatives in the final dataset is 2,771. Table 1 provides an overview of the distribution of derivatives among centuries.

	N*-ity*	**N*-ness***	**N$_{overall}$**
20th century:	344 (61%)	220 (39%)	564 (100%)
19th century:	733 (49%)	759 (51%)	1,492 (100%)
18th century:	306 (43%)	408 (57%)	714 (100%)

Table 1 *OED neologisms - general overview*

Two things are particularly noteworthy. The first thing is the disproportionately high number of neologisms from the 19th century. The differences between the 19th century on the one hand and the 18th and 20th centuries on the other hand are fully in line with the general distribution of first attestations in the OED, which can, for example, be examined on the OED *Timelines* page (http://www.oed.com/timelines, accessed on June 18, 2013). The second thing that is noteworthy about Table 1 is that the type frequencies attested in the OED clearly do not reflect the greater general productivity of *-ness* as documented in the literature. Not only are *-ity* and *-ness* derivatives in general represented by quite similar type frequencies; there is even a consistent increase of *-ity*'s share in the distribution from the 18th century onwards.

It is clear that type frequencies gleaned from dictionary data are not a very good measure of overall productivity (cf. e.g. Baayen & Renouf 1996, cf. Plag 1999: 98 for quantitative evidence that *-ness* words are not so likely to make it into a dictionary).[6]

[6] On a more general level, a potential problem with using dictionary data to investigate productivity is that dictionary data provide a somewhat biassed record of the range of functions that a derivational category may have. In particular, functions that are discourse referential rather than conceptual are unlikely to be represented adequately in the dictionary (cf. Kastovsky 1986, Baayen & Neijd 1997 for a motivation of the distinction and discussion). It will be an interesting matter for future research to see

Nevertheless, the analysis to follow will show that the proportionate distribution of the two suffixes undergoes consistent changes over the three centuries under scrutiny. This will be interpreted as indicative of changes in the productivity profiles of the two suffixes.

In order to test the claim that *-ness* is less restrictive than *-ity* in terms of the syntactic category of its bases, all data were coded manually for the syntactic category (word class) of the base. Pertinent categories were taken from the OED classification of the base of the *-ity* or *-ness* derivative, with the addition of two minor categories. These are 'bound form' and 'phrase'. The category 'bound form' was used in cases in which, according to the OED, the base of the nominal derivative is not attested as an independent word in English. There are two types of bases where this is the case: The first type comprises bound roots as, for example, in the word *iracundity*. The second type comprises bases involving adjectival suffixes which only occur in the nominal derivative. An example is the word *randomicity*, where the base ?*randomic* (ending in the adjectival suffix *-ic*) is unattested.[7] The category 'phrase' comprises combinations of words that are not listed by the OED to be words of English. The list of coded syntactic categories is given in Table 2.

syntactic category of the base	**example words (century of first attestation)**
adjective	*lovability* (19th century), *nerdishness* (20th century)
noun	*perspectivity* (20th century), *moneyness* (20th century)
verb	*relaxity* (18th century), *oughtness* (19th century)
adverb	*onceness* (19th century)
preposition	*betweenity* (18th century), *betweenness* (19th century)
pronoun	*I-ness* (19th century)

whether and in how far the results obtained in the present study carry over to discourse referential usages of *-ity* and *-ness*, for example in a corpus. Thanks to an anonymous referee of this article for pointing this out!

[7] Note, furthermore, that cases which have traditionally been classified as cases of systematic stem truncation were not classified as bound forms. Thus, I assume that the bases for words like *homogeneity*, *miscellaneity*, *anonymity*, and *incontinuity* are *homogeneous*, *miscellaneous*, *anonymous*, and *incontinuous*, respectively.

wh-pronoun	*whenness* (20th century)
particle	*notness* (20th century)
bound form	*iracundity* (19th century), *arity* (20th century, based on the suffix *-ar*)
phrase	*know-nothingness* (19th century)

Table 2 *Coding: syntactic category of the base*

In order to investigate the influence of morphological category information on deadjectival formations, adjectival bases were coded for the pertinent adjectival suffix or, in cases in which there was no suffix, for its absence. The coding will henceforth be referred to as the 'morphological category' of the adjectival base. The definition of morphological complexity that underlies the coding was as generous as possible. A base was coded as containing an adjectival suffix if (a) its syntactic category is an adjective, and (b) the form of the suffix is known to be a common adjectival suffix in English. This meant, then, that roots could be bound (as, for example, in words like *arable* ♦ *arability*, *comical* ♦ *comicality*, *didactic* ♦ *didacticity*; cf. Bauer et al. 2013: chpt. 3.2 for discussion).

Finally, the data were broadly coded in terms of the transparency of their first attested meaning. The purpose of that coding was to test whether *-ity* and *-ness* can be used as synonyms in the sense that they both denote transparent properties. This is important in order to assess claims in the literature that avoidance of synonymy plays a role in the competition between the two nominalisers. This idea is a crucial aspect of the framework proposed by Lindsay and Aronoff (2013), Lindsay (2012), and an implication of the account by Riddle (1985). 'Transparency' in the present analysis means that the derivative encodes a property or state derived from the meaning of the base word. Technically, the coding was done semi-automatically on the basis of the description of its first attested meaning in the OED (cf. Hay 2001, Hay 2003, who uses a very similar methodology). If the meaning description contained the words 'property', 'quality', 'state', or 'condition', the derivative was coded to be 'clearly transparent'. In all other cases the transparency of the derivative was coded to be 'unclear'. Examples of

meaning descriptions for the two values of the variable 'transparency' are given in (2). All data in (2) are 20th century neologisms.

(2) (a) descriptions indicating clearly transparent items

aggregability	The property or condition of being aggregable, the tendency to aggregate or be aggregated; the degree to which things, esp platelets or other cells, aggregate
peakiness	The state or condition of being peaky (peaky); sickly or undernourished quality or appearance
scientificity	The quality of being scientific; scientific character

(b) descriptions indicating unclear transparency

aeroelasticity	The branch of mechanics dealing with the interaction of aerodynamic, elastic, and inertial forces, esp in aircraft structures
affectlessness	Detachment alienation; incapacity to feel emotion
bibulosity	Addiction to tippling

Whereas we can safely assume that all data coded as 'clearly transparent' are indeed transparent in the above-described sense, it is clear that the group of items like those in (2b) constitutes a heterogeneous group with respect to their transparency. Quantitative analysis of these data can therefore only be used to compare the proportion of clearly transparent items among *-ity* and *-ness* neologisms. Crucially, however, they cannot be used to estimate the overall degree of transparency of each suffix.

3.2 *Productivity of -ity and -ness in 20th century English*

This section will investigate in how far the coded factors account for significant differences in the distribution of *-ity* and *-ness*. We will see that the distribution of the two nominalisers is fully complementary for some factors whereas other factors can be used to describe only relative preference patterns. Yet other factors turn out not to be convincingly supported by the data. In what follows each factor will be discussed in turn.

Table 3 provides an overview of the distribution in terms of the syntactic category of the bases. Due to the small number of occurrences, the syntactic categories 'adverb', 'preposition', and 'pronoun' are summarised as 'minor categories'.

syntactic category of the base	***-ity* derivatives**		***-ness* derivatives**	
	N	**%**	**N**	**%**
adjective	326	94.8%	186	84.5%
noun	7	2.0%	14	6.4%
bound form	11	3.2%	0	
phrase	0		10	4.5%
minor categories (adverb, preposition, pronoun)	0		10	4.5%
Total	344	100.0%	220	100.0%

Table 3 *-ity / -ness derivatives by syntactic category of the base, 20th century*

Table 3 shows that the distribution of *-ity* and *-ness* is complementary for bound and phrasal bases as well as the group labelled 'minor catgories'. For nominal and adjectival bases we find both *-ity* and *-ness* words attested. As expected, both nominalisers occur predominantly with adjectival bases. Differences between the proportions of nominal bases among *-ity* and *-ness* derivatives are statistically significant (Yate's $\chi^2 = 5.858$, df = 2, $p < 0.05$), but subject to the obvious caveat that pertinent numbers are rather small. Examples of *-ity* and *-ness* derivatives with non-adjectival bases are provided in (3) and (4).

(3) non-adjectival bases of *-ness* derivatives

phrases	**nouns**	**other bases (exhaustive list)**
not-thereness	nannyness	aboutness
on-the-makeness	noodleness	againstness
as-suchness	nounness	beyondness
middle-of-the-roadness	partyness	meness
		notness
		usness
		weness
		whenceness
		whoness
		withness

(4) non-adjectival bases of *-ity* derivatives

bound bases	**nouns**
helicity	orgonity
randomicity	chlorinity
zygosity	perspectivity
arity	picayunity

The focus of discussion will now be restricted to derivatives with adjectival bases, to test the claim frequently made in the literature that the choice between *-ity* and *-ness* is influenced by the morphological category of adjectival bases. Table 4 provides a first overview of how these morphological categories are represented in the database. Note that only those classes are listed which are represented by at least 10 derivatives in the database (N = 545).

morphological class	**N**	**%**	**examples**
-able	109	21.54%	achievability
-al	78	15.42%	directionality
-ic	55	10.87%	iconicity
-y	55	10.87%	nerdiness
-ive	51	10.08%	destructivity, informativeness
no suffix	42	8.30%	multiplexity, alienness
-ed	29	5.73%	handedness
-ous	24	4.74%	bibulosity, murmorousness
-ing	12	2.37%	terrifyingness
-ar	12	2.37%	modularity
-ish	10	1.98%	kittenishness
-less	10	1.98%	affectlessness
Total	506	100.00%	

Table 4 *Morphological classes of bases, 20th century*

The vast majority of bases (464 of 506 bases, i.e. 91.70%) contain adjectival suffixes. Furthermore, not all adjectival suffixes are equally well represented in the database, with *-able* clearly forming the largest group. Also note that Germanic affixes are not well represented – we only find *-y* and *-ed* in numbers that are somewhat substantial.[8]

[8] The finding that *-ness* is especially frequent with adjectival bases ending in *-y* and not very frequent with other bases ending in Germanic suffixes is fully in line with what Baayen and Renouf (1996) find in their corpus data.

The mosaic plot in Figure 1 shows how *-ity* and *-ness* are distributed for bases with an adjectival suffix. Again, only morphological categories with a minimum of 10 attestations are taken into account.

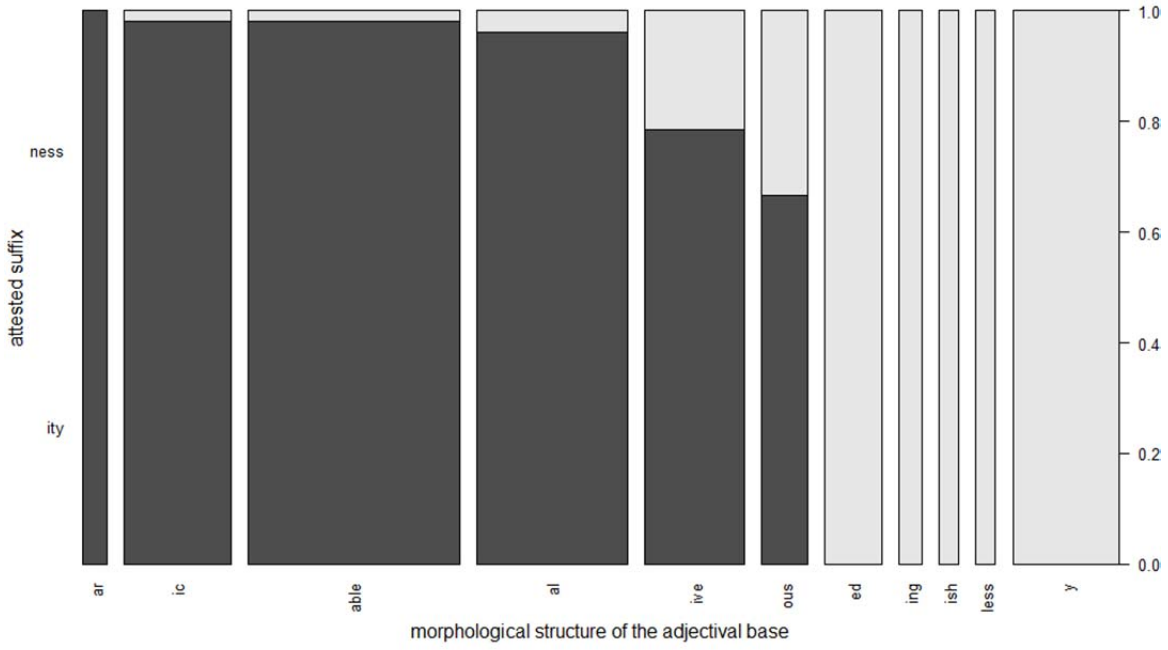

Figure 1 *20th century: attested distribution of -ity and -ness by morphological category of the base, N = 445*

The bar plot represents the proportion of *-ity* and *-ness* derivatives as dark and light grey areas respectively, which are scaled to 100% on the y-axis. Each bar corresponds to a morphological category (labelled on the x-axis), and the width of the bars is proportionate to the pertinent number of observations.

As is to be expected, *-ness* is the only nominaliser that occurs on bases ending in *-ed*, *-ing*, *-ish*, *-less*, and *-y*. With the exception perhaps of *-ar* bases, *-ity* never occurs categorically in any of the remaining classes. However, it is very strongly preferred to *-ness* with bases ending in *-ic*, *-able*, and *-al*. *-ity* also clearly constitutes the most frequent nominaliser for adjectival bases ending in *-ive* and *-ous*. The very strong preference of *-ity* for bases with the abovementioned suffixes provides a challenge to theories that explain rivalry between *-ity* and *-ness* in terms of overall differences in productivity. The *-ity* preferences found for some of the morphological categories in the

OED data receive independent support from the findings of recent corpus-based work (Lindsay 2012, Lindsay & Aronoff to appear), which has documented a clear *-ity* preference for the suffixes *-al*, *-able*, and *-ar* on the basis of the distribution of token frequencies among rival word pairs.

The final factor to be looked at in this section is the semantics. Figure 2 plots the proportion of clearly transparent items in the 20th century dataset.

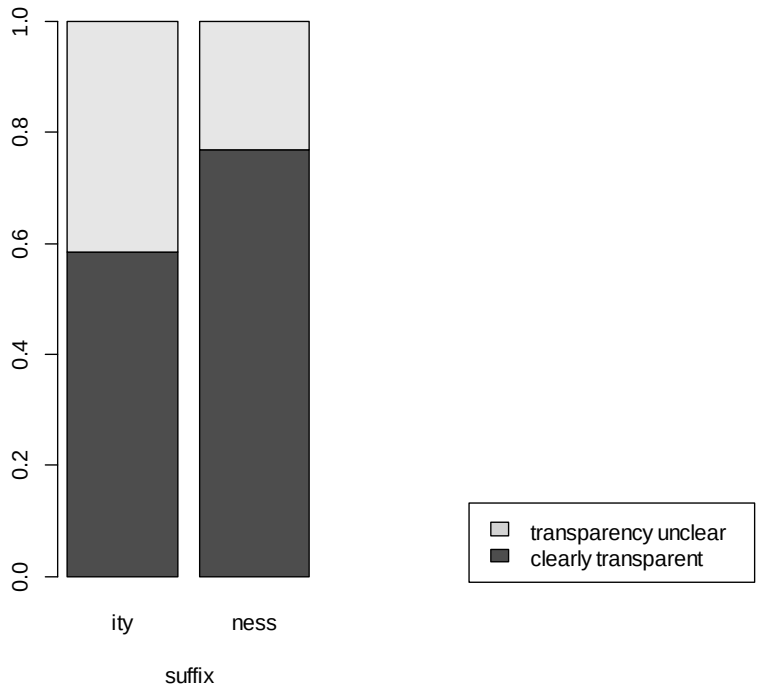

Figure 2 *Transparency of -ity and -ness derivatives, 20th century (N = 564)*

The clear majority of both *-ity* and *-ness* derivatives in the dataset denote a property or state derived from the base word (*-ity*: 58.5%, *-ness*: 76.8%). This shows that *-ity* and *-ness* are productively used as synonymous suffixes. The second observation is that the proportion of clearly transparent items is significantly higher for *-ness* than for *-ity* (Yates' χ^2 = 17.07, df = 1, p = 3.611e-05). This provides support for the view that there are differences between *-ity* and *-ness* which can broadly be characterised as semantic. What these differences might be, however, is still far from clear. Given the very crude and superficial coding of transparency here, it is a matter of future research to investigate in more detail the motivation of this difference.

The following section will trace the development of the distribution of *-ity* and *-ness* in the 20th century neologisms by looking at the corresponding distributions of neologisms from the 18th and the 19th centuries.

3.3 *Diachronic development, 1700 - 2000*

The main thing to be learned from a comparison of the variation between *-ity* and *-ness* neologisms found in the 20th century with that in the 18th and 19th centuries is that the 20th century patterns constitute only one step in a consistent development that has taken place during the three centuries. This development seems to be form-based rather than meaning-based. In particular, it is most visible when we look at the distribution of nominalisers among adjectival bases belonging to different morphological categories. Recall, however, that the morphological category of the base is used here as a convenient label to refer to similarity relationships between words which may in fact be purely form-based. No commitment is made here to the theoretical claim that it is indeed the morphological category that is decisive (cf. section 2 above for discussion).

Figure 3 provides an overview of how morphological categories are distributed among adjectival bases in the three datasets, plotting their overall frequencies for each of the three centuries under investigation. For reasons of readability, separate plots are drawn for bases with Latinate adjectival suffixes (upper left panel), Germanic adjectival suffixes (upper right panel), and for unsuffixed bases (lower left panel). Note that different scales are used in the y-axes of the three plots. For the 18th and 19th century data, the numbers underlying the plot are provided in the appendix to this paper. Pertinent numbers for the 20th century data can be found in Table 4 above.

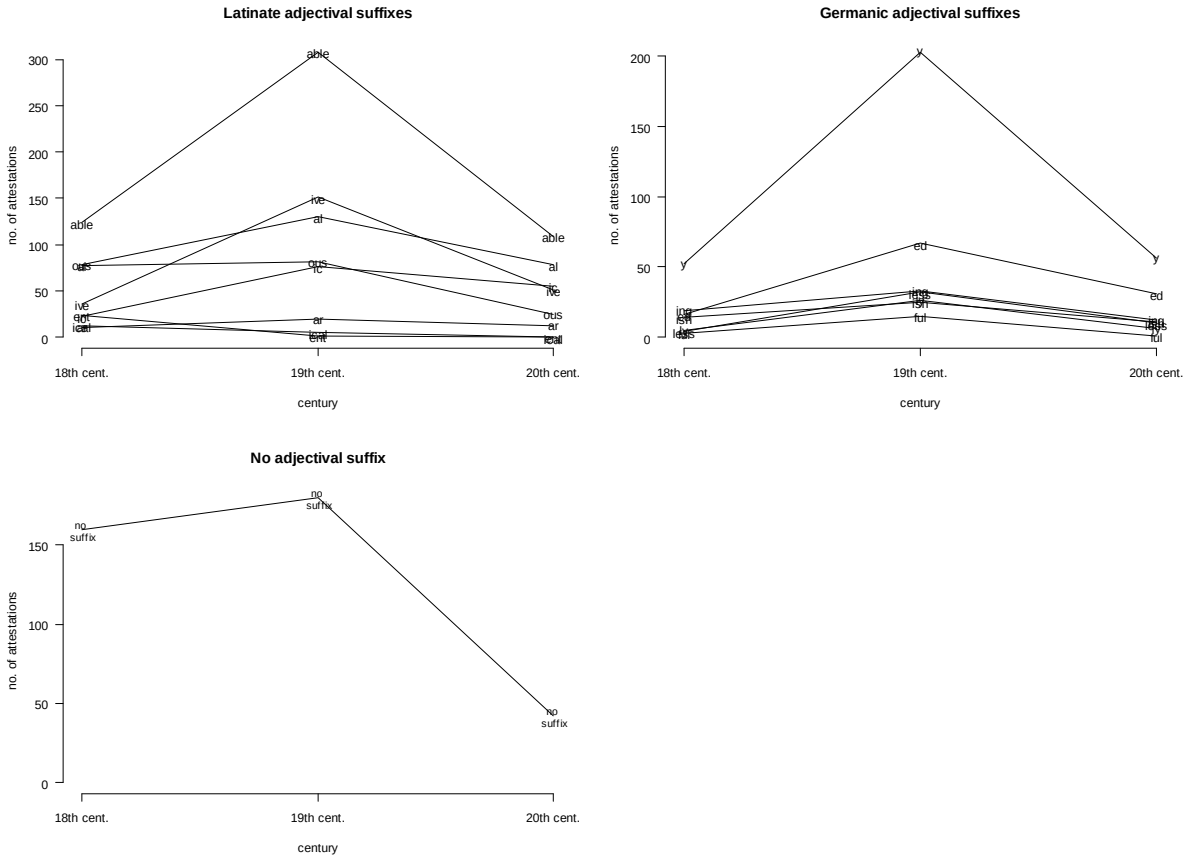

Figure 3 *Morphological categories of adjectival bases, diachronically ($N_{18th\ cent.}$ = 673, $N_{19th\ cent.}$ = 1,379; $N_{20th\ cent.}$ = 512)*

Adjectives ending in *-able* form the largest group of complex adjectival bases in all centuries. Furthermore, we see that the disproportionate size of the 19th century dataset is reflected as peaks in the distribution especially of the frequent base categories. Interesting differences between the different centuries can be found especially for the relative frequency of unsuffixed adjectival bases, which are strikingly frequent in the 18th century data if compared to the more recent data. Among Latinate adjectival suffixes, we see a decline in frequency of *-ous*, and a rise especially of *-ic*.

Comparing the distributions of *-ity* and *-ness* for the morphological categories of bases, we see that bases with frequent Latinate suffixes show a consistently increasing tendency to take *-ity* rather than *-ness*. By contrast, no variation is observed for *-ness*.[9] Figure 4 provides a pertinent overview for the Latinate suffixes. Note that it only

[9] A potential exception to the generalisation that the Germanic suffixes consistently take *-ness* is the item *forlornity*, first attested in the 19th century. However, the morphological status of the base *forlorn* as a participle form is unclear.

incorporates classes with a minimum of 10 attestations for each century.

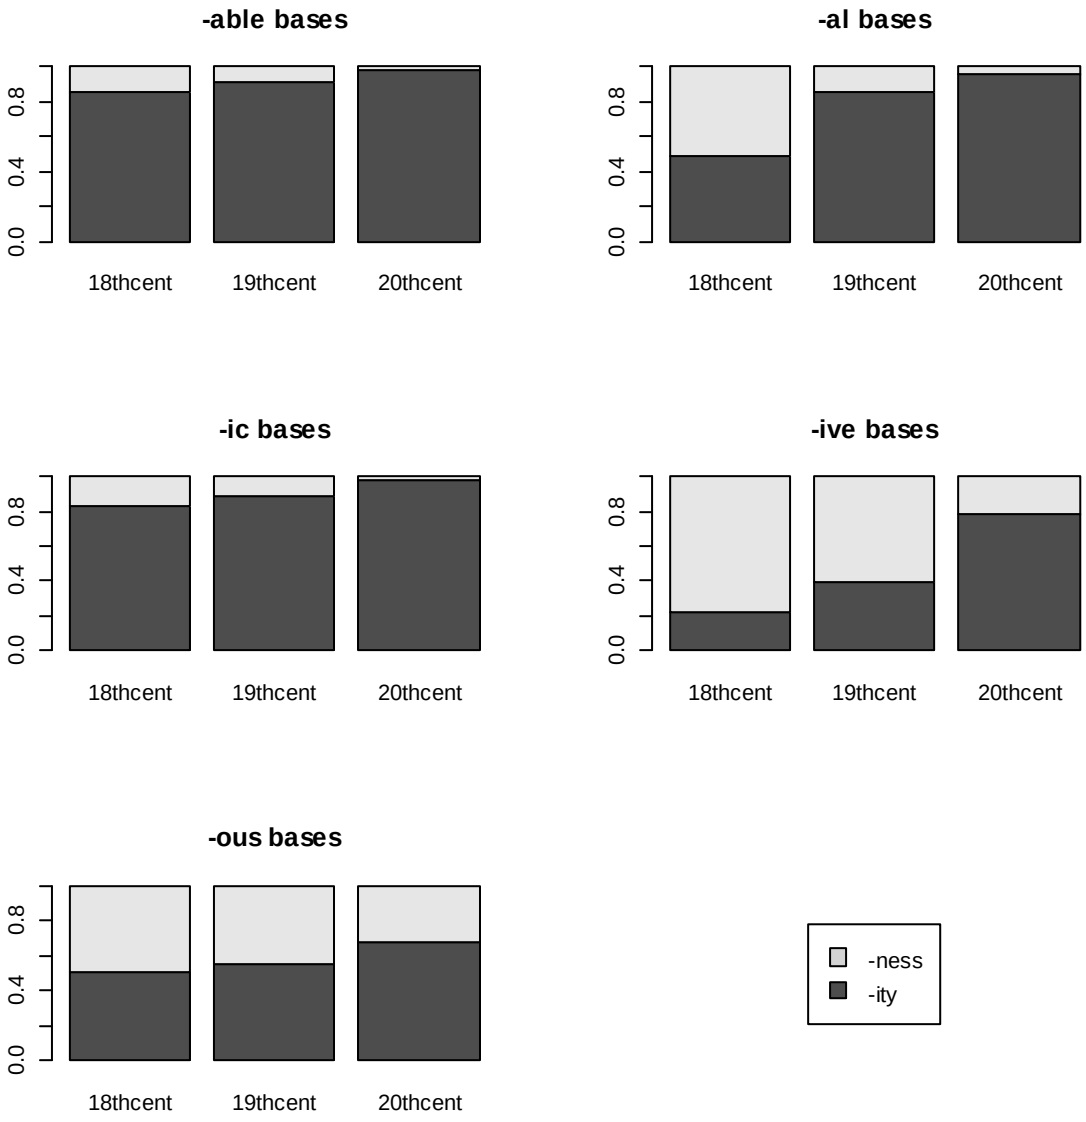

Figure 4 *The distribution of -ity and -ness neologisms among frequent Latinate categories, diachronically (N = 1,425)*

For all Latinate suffixes, Figure 3 shows a steady and consistent increase of *-ity's* share in the distribution over the last three centuries, so that in the 20th century data, there is a very clear majority for *-ity* in all pertinent morphological categories. Furthermore, note that morphological classes differ in terms of the degree of their *-ity* preference. Whereas *-ity* is almost categorical for *-able*, *-al*, and *-ic* bases in the 20th century data,

variation is considerably larger among *-ive* and *-ous* bases. This variation is contrasted with the 18th century, where *-ity* forms already constitute a clear majority among *-able* and *-al* bases, but where *-ness* is still the clear majority nominaliser for *-ive* bases and choice of *-ity* or *-ness* seems to be more or less random among *-ic* and *-ous* bases. The tendencies observed here are difficult to reconcile with the idea that the development of the division of labour between *-ity* and *-ness* is a reflex of developments that have been triggered by whole-word borrowings in the Middle English period. Especially for *-al*, *-ive*, and *-ous* the trend towards *-ity* seems to be a development that has taken place much later.

With respect to the syntactic category of the base, not much seems to have changed in the last three centuries in terms of how *-ity* and *-ness* derivatives are distributed among different types of bases. What has, however, changed to some extent is the overall distribution of base categories. The latter change is illustrated in Figure 5, where adjectival bases and non-adjectival bases are plotted in two different panels for convenience. A table containing the pertinent numbers can again be found in the appendix of the paper.

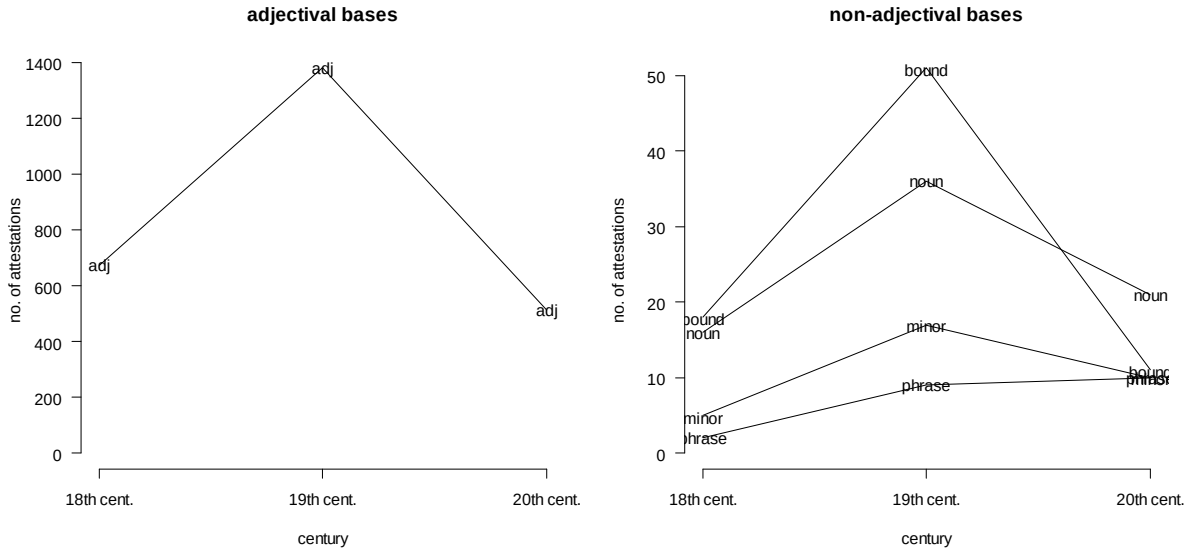

Figure 5 *Syntactic categories of bases, diachronically ($N_{18th\ cent.}$ = 714, $N_{19th\ cent.}$ = 1,492; $N_{20th\ cent.}$ = 564)*

We see that the absolute number of adjectival bases represented in the OED data drops from 673 in the 18th century to 512 in the 20th century. The proportion of adjectival

bases in the overall set of bases thus drops from 94.3% in the 18th century dataset to 90.8% of the 564 bases in the 20th century dataset. This drop in adjectival bases is compensated for by a rise in number of derivatives that are based on nouns, phrases, and the categories labelled 'minor' in this paper (verb, adverb, preposition, pronoun, wh-pronoun, and particle).

Figure 6 provides an overview of the diachronic distribution of *-ity* and *-ness* for these base categories.

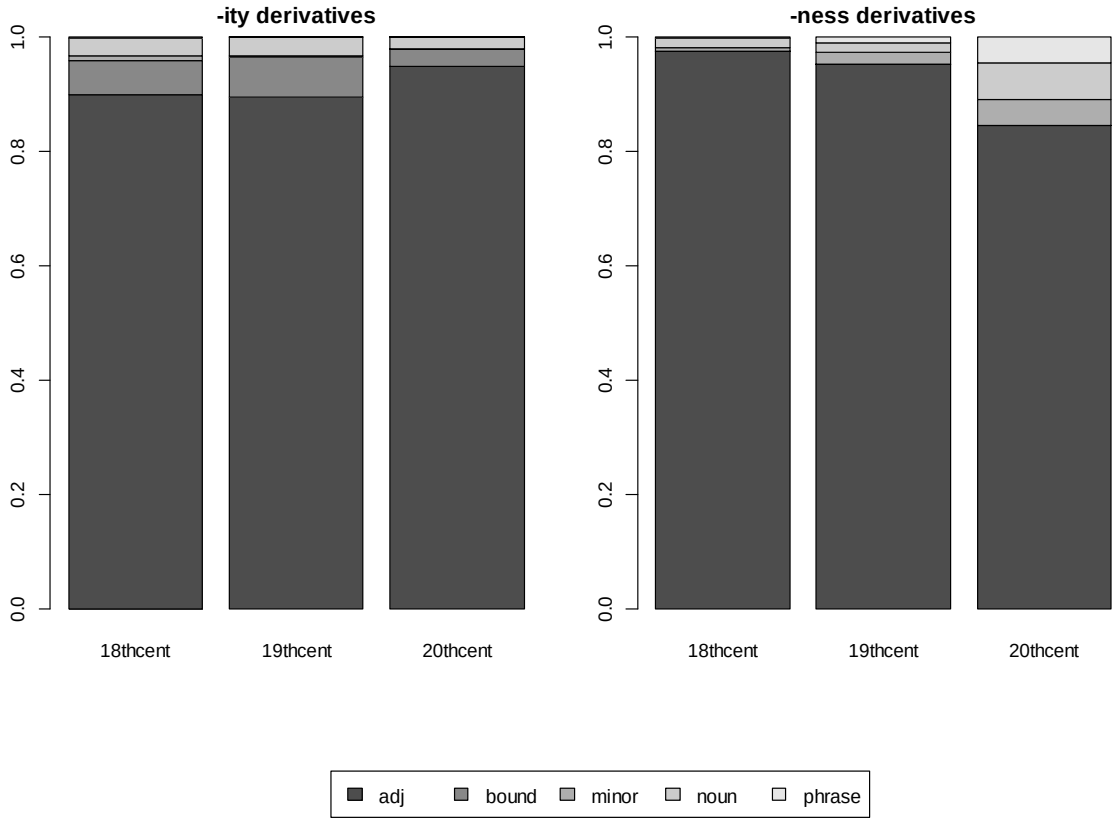

Figure 6 *Syntactic category of bases, diachronically (N = 2,770)*

Figure 6 suggests that unlike *-ity*, *-ness* has diversified in terms of its potential bases; in particular, it has increasingly been attested as a nominaliser for phrases, nouns, and the categories listed as minor categories.[10] By contrast, the decrease in the proportion of non-adjectival bases of *-ity* derivatives is mainly due to a decrease in the share of bound bases in the distribution. Note, however, that, even if pertinent numbers are small, also

[10] Note that there is the possibility that the diversification that can be observed here is an artifact of lexicographic policies in the OED.

bases labelled 'minor categories' were more frequent with *-ity* in the historical than in the contemporary data. Particularly convincing examples are the items from the 18th century, which show that *-ity* could be attached to prepositions (*betweenity*) and verbs (*relaxity*, *flexity*).

The final issue to be addressed in this section concerns semantic transparency. Figure 7 plots the proportion of clearly transparent items for all three centuries. The figure includes Figure 2 for convenience.

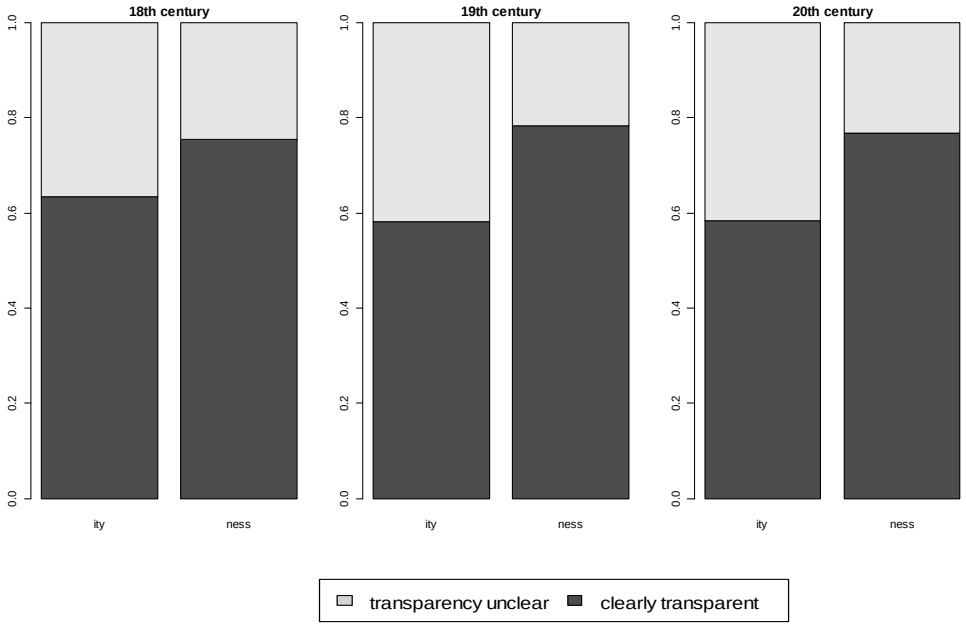

Figure 7 *Transparency of -ity and -ness derivatives, diachronically (N = 2,775)*[11]

The graphs illustrate the fact that distributions of *-ity* and *-ness* neologisms from the three centuries are highly similar in terms of the proportion of clearly transparent items in the sense defined in this paper. The proportion of clearly transparent items among *-ity* neologisms is 63.60% for the 18th century, 58.23% for the 19th century, and 58.49% for the 20th century. The proportion of transparent items among *-ness* neologisms is 74.75% for the 18th century, 78.28% for the 19th century, and 76.80% for the 20th century.

In sum, the analysis of the diachronic data in this section has provided evidence

[11] Note that N is slightly bigger (2,775) than in the rest of the section where N was 2,771. The reason is that in the 19th century data, there are four items for which the OED lists two different meanings as first attested meanings. These are *correality*, *viability*, *brashness*, and *peakedness*. Each meaning is given a separate count here, whereas counts in the rest of the paper are form-based.

for a developmental dynamics in the division of labour between *-ity* and *-ness*. Crucially, this dynamics does not support an explanation in terms of a simple regularisation of the system. Thus, *-ity* develops to become the dominant nominaliser for adjectival bases ending in *-able*, *-al*, *-ic*, and *-ive* (and maybe *-ous*). However, the development cannot be explained in all cases as the majority nominaliser taking over the morphologically defined class of bases. A case in point are bases ending in *-ive* and *-al*, for which the majority nominaliser is *-ness* in the 18th century, and where we see a subsequent increase in variation and a decrease in the predictability of nominaliser selection. With respect to the syntactic categories of potential bases, we have seen the converse situation. Thus, *-ity* seems to be more restricted in terms of its bases nowadays as compared to the 18th century. *-ness*, by contrast, seems to have undergone the opposite development. Bases of *-ness* suffixation are more diverse in the 20th century than they were in the 18th century. Again, the evidence speaks against regularisation as the motor of the development because such an assumption could not explain why bases of *-ness* suffixation become more unpredictable over time.

3.4 *Section summary*

Two important things emerge from the findings in this section. The first is that the OED data are a valid resource for the analysis of the division of labour between *-ity* and *-ness*. The 20th century distribution is fully in line with what the few available corpus studies show about the domain-specificity of the two nominalisers. Thus, Baayen and Renouf (1996) as well as Lindsay (2012) find that *-ity* is more productive than *-ness* in some domains involving Latinate suffixes. However, the present analysis suggests that the productivity difference may be even more pronounced than hitherto assumed. Furthermore, the distribution of *-ness* confirms earlier corpus-based findings in Baayen and Renouf, who show that *-ness* is particularly productive with adjectival bases ending in *-y*, and much less productive with bases ending in other Germanic suffixes such as *-ed* and *-ing*. The close correspondence of findings between the present study and that corpus study show that *-ness* derivatives are reasonably well represented in the OED data to allow for a comparison with *-ity*.

The data on diachronic development are fully in line with Riddle's claim that *-ity* formations tend to be less likely to have a transparent property meaning than *-ness* formations. However, we saw that the clear majority of both groups of nominalisations

have clearly transparent meanings in both the synchronic and the diachronic data. This indicates that even though meaning differences may be a factor in the competition between the two nominalisers, the importance of this factor is a relative one. Furthermore the diachronic data are in line with what Lindsay (Lindsay 2012, Lindsay & Aronoff to appear) hypothesises about the development of niche productivity in general. Again, however, the development may in fact be more pronounced than hitherto noted. Given that the data reflect the more recent history of the language, they show that the process has not lost its historical dynamics after Middle English. This is particularly true for the behaviour of adjectival bases ending in *-al* and *-ive*, where we observe a drastic change in their preferences. Also, the results presented here suggest that it is unlikely that the developments are still the consequence of the developments in Middle English. Instead, more research is needed to establish whether the development is due to factors that have arisen much later.

4 Introducing an analogical approach to -ity / -ness rivalry

In what follows I will propose and test an analogical model of *-ity* and *-ness* selection. The task of the present section is to provide a general introduction to the basic mechanisms of the model. The details that are necessary to understand synchronic productivity patterns and diachronic development in rivalry will then be discussed in connection with the concrete findings of the present model in section 5. Note that the description here will not incorporate technical details but focus on a rather broad, conceptual level. For the mathematical and technical details the interested reader is referred to pertinent descriptions in the literature, especially those in Skousen (2002a, 2002b, 2005, 2009).

The particular framework used in the present study (Analogical Modeling, Skousen 1989, Skousen 1992, Skousen et al. 2002 et seq.) constitutes an instantiation of what I will henceforth refer to as an analogy-based approach to linguistic categorisation (or 'analogical approach', for short). The claim on which the approach is based is by no means new. For word-formation, the assumption is that new words are formed on the basis of existing similar words. More specifically, new complex words are formed from their bases on the basis of existing base-derivative pairs in the mental lexicon (cf. e.g. Becker 1990, Bauer 2001: 75–96, Anttila 2003, Arndt-Lappe in press for an overview

of the basic tenets of analogical approaches to word-formation). The crucial characteristic here is that analogy happens online in linguistic tasks. This view is different from many current, especially generative theories in that such theories often place an emphasis on the rule as the central mechanism of linguistic generalisation. Rules, in turn, are abstractions that are assumed to operate independently of individual exemplars in the lexicon in such approaches.

The term 'analogy-based approach' can thus be seen as a cover term for quite a variety of different approaches which are united by the view that linguistic generalisation happens online on the basis of the lexicon. The most prominent among these are exemplar-based theories (cf. overviews in e.g. Gahl & Yu 2006, Wedel 2006), computational analogical approaches of different types (in morphology, especially: AM, Skousen 1989, Skousen 1992, Skousen et al. 2002 et seq.; TiMBL: Daelemans et al. 1999 et seq., Daelemans & van den Bosch 2005, 2013; Minimal Generalization Learner: Albright & Hayes 2003, Albright 2009), and, to some extent, construction-based theories (esp. Booij 2010). Note that the term 'exemplar-based approaches' is often also applied to the computational approaches cited above. In this paper, the term 'analogical approach' is preferred because in recent linguistic work much discussion of exemplar-based models is centered around the role of phonetic detail in exemplar representation. By contrast, in work applying computational analogical models to morphological phenomena, there is a tendency to focus on more abstract components of lexical representations. In this sense, then, the term 'analogical approaches' is chosen in order not to invite any biasses with regard to the structure of representations.

In the present model, *-ity* and *-ness* selection is viewed as a classification task. For each new base word that is given to the algorithm as a test item, the algorithm decides between two different outcomes, 'ity' and 'ness', on the basis of the distribution of similar items bearing either *-ity* or *-ness* in the lexicon. Similarity is computed on the basis of the set of coded features that is assigned to each exemplar in the lexicon by the researcher. In our model, two types of information will be used: information on the form of the base and syntactic information. The form information consists of six features which encode the phonemic makeup of the two last syllables of the base. Information on the morphological category of the base is therefore implicit in this coding, which allows for a more finegrained comparison of the formal properties of

exemplars than a comparison that is based on membership in a morphological category. The coding of syntactic information covers those syntactic categories that have emerged in section 3.2 to be most relevant for *-ity* and *-ness* selection. The classification encodes whether the base is a word in English (collapsing the categories 'adjective', 'noun', and 'minor' from section 3 into one umbrella category) or whether it is a bound form or a phrase.

Figure 8 provides a graphical illustration of the general architecture, using the base *perspective* as an example. For *perspective*, the OED lists *perspectivity* as a 20th century neologism. The lexicon represented in Figure 8 constitutes a heavily simplified version of the lexicon that will be used in the synchronic simulation study in section 5.

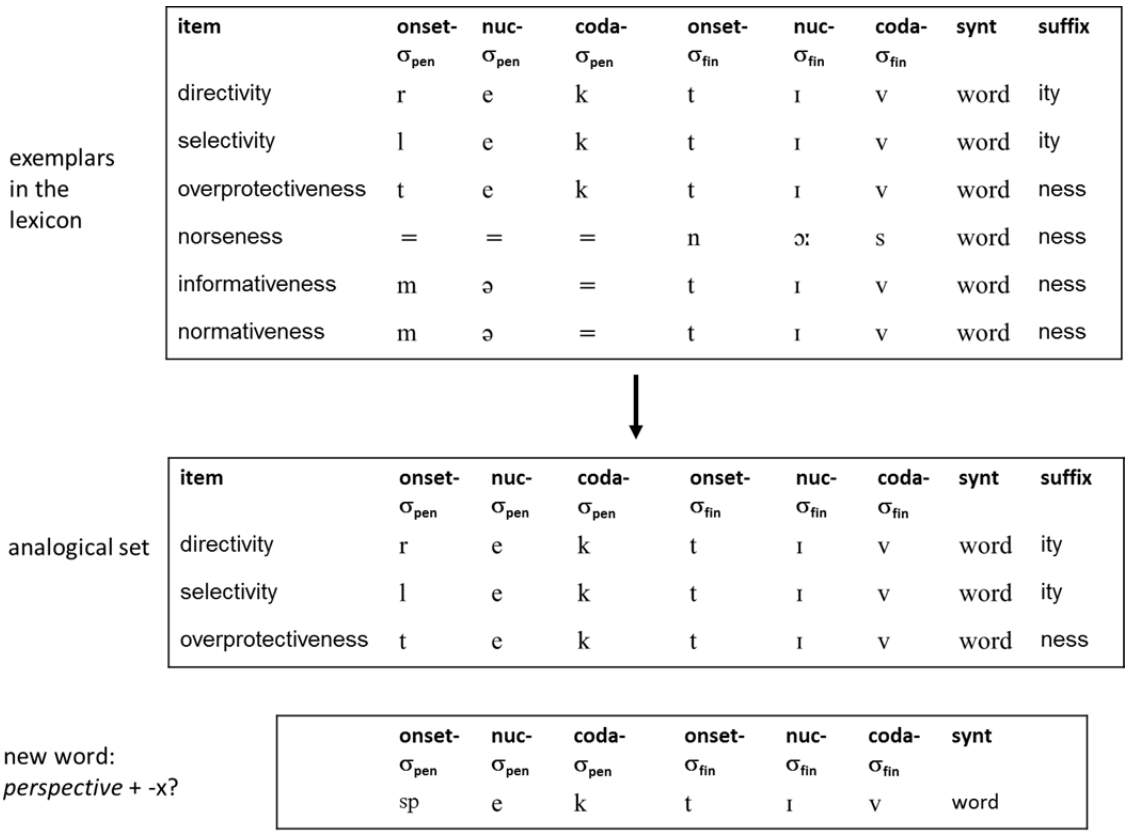

Figure 8 *The general architecture of an analogical model*

The central component of the system is the lexicon, i.e. the set of stored exemplars. In the lexicon each item is represented as a structured set of coded features. In our case, the set involves eight features, six of which represent the phonemic makeup of the last two syllables of the base. Each feature represents a syllabic constituent (onset, nucleus, and coda), and the symbol '=' indicates that the position is not filled by a segment (similar

representations have been employed frequently in similar tasks, e.g., in Eddington 2002, Chapman & Skousen 2005, Daelemans & van den Bosch 2005, Keuleers 2008, Chandler 2010, van den Bosch & Daelemans 2013). Another feature encodes the syntactic category of the base; this coding broadly distinguishes words (i.e. free forms) from bound forms and phrases. In Figure 5, all items in the lexicon are derivatives based on words. The last feature in the set specifies whether the attested nominaliser is *-ity* or *-ness*.

If a new word is to be classified, the system extracts from the lexicon a group of exemplars which are similar to the new word. This group is known in the AM literature as the analogical set. The exemplars in the analogical set serve as analogues on the basis of which the nominal suffix will be assigned to the new word. In the example in Figure 8, all exemplars in the analogical set share as many as six of the seven coded features with the new word (i.e. all items end in /ektɪv/ and are attested as words in English) and are thus all equally similar to the target. An obvious question is how the system determines which exemplars end up in the analogical set. For example, the small lexicon in Figure 8 also contains two exemplars, *informativeness* and *normativeness*, which are reasonably similar to the new word (sharing three phonological features and the word class with the target), but are not included in the analogical set. The crucial feature that distinguishes AM from many other pertinent models is that the degree of similarity that is relevant for exemplars to be included in the analogical set is decided for each new word individually. The rationale that underlies the procedure is that while the model will always incorporate maximally similar items, items with lower degrees of similarity will be incorporated only if that incorporation does not lead to greater uncertainty with respect to the classification task.[12] As we will see in section 5.2 below, this is relevant for an explanation of how the different productivity profiles of *-ity* and *-ness* emerge from the model.

On the basis of the analogical set, the nominalising suffix for the new word is predicted. This prediction is probabilistic, reflecting the distribution of nominalising suffixes in the analogical set. Specifically, the algorithm gives the probability of each

[12] Technically, the procedure involves a systematic comparison of all possible constellations of features (termed 'contexts') shared between the target item and the exemplars in the lexicon, which gives privilege to constellations where more features are shared over those where fewer features are shared. Cf. especially Skousen (1992, 2002a, 2002b, 2005) for discussion and explanation.

possible value being assigned, based on the distribution of these values amongst exemplars in the analogical set. In calculating these probabilities, AM takes into account the degree of similarity between an exemplar and the new item, as well as the number of exemplars with a particular set of features. The more similar an exemplar is to the new item, the more weight it receives, and the more exemplars that share a particular set of features, the greater the weight assigned to each of them. The latter fact is particularly important in cases where a feature constellation is particularly frequent in the lexicon. In a classification task where exemplars sharing that feature constellation are part of the analogical set, they may outweigh exemplars that are more similar to the target word but that are smaller in number.

In order to better assess the overall performance of the model, probabilistic outcome predictions are often translated into categorical predictions for each new word. This means that probabilities of a given outcome that are greater than 50% are taken to be a categorical prediction for that outcome to occur. This procedure is well-established in the literature applying analogical models to linguistic tasks (cf. e.g. papers in Skousen et al. 2002, Daelemans & van Bosch 2005, for examples and discussion). In the discussion of the simulation study in this paper reference will be made sometimes to the probabilistic predictions and sometimes to the transformed categorical predictions.

5 THE SIMULATIONS

In what follows I will report on a series of simulation experiments with AM that use the OED data to test how the analogical model accounts for the patterns and developments observed in the analysis in section 3. After a brief introduction to the general setup of the simulations in section 5.1, section 5.2 will provide an overview of the overall predictive power reached in the simulations. Section 5.3 will then explain how the AM simulation accounts for the productivity profiles of *-ity* and *-ness* in the 20th century simulation. Section 5.4 will focus on the modelling of diachronic change from the 18th to the 20th century.

5.1 *Coding and general setup*

In all simulation experiments, AM was set the task of predicting the nominalising suffix for bases of *-ity* and *-ness* derivatives from the OED dataset. As a lexicon, the algorithm

used a corresponding set of *-ity* and *-ness* derivatives from the OED dataset. The basic shape of the coding has already been exemplified in Figure 8. The phonological coding was based on the machine-readable transcription which can be accessed in the OED Online version by clicking on the IPA transcription. The transcriptions were then edited manually to reduce inconsistencies. It is important to note that alignment of the phonological features was position-specific. This means that onsets, nuclei, and codas of the penultimate and the final syllable were always compared only with corresponding onsets, nuclei, and codas from the same syllable. Furthermore, the symbol '=', which indicates that the relevant syllabic position is not occupied by a segment (cf. Figure 8 for illustration), was treated as a valid feature value. The effect was that the model was able to incorporate into the computation of similarity some effects that take place on a higher level of structural organisation than the phonemic makeup. One constellation that produced a noteworthy effect here was the absence of a coda in the penultimate syllable of the base, another was the constellation with three '=' values for the features of the penultimate syllable, which indicated that the item is monosyllabic.

For the coding of morphosyntactic category, the coding introduced in section 3.1 was transformed into broader categories. In particular the categories 'adjective', 'noun', 'verb', 'adverb', 'preposition', 'pronoun', 'wh-pronoun', and 'particle' were conflated to create a general category labelled 'word'. This category then comprises all bases which are minimal free forms in English. Note that the reasons for the conflation of word categories were purely practical in nature in that, in terms of the predictive power of the model, nothing could be gained from a more finegrained distinction. This is almost certainly due to the fact that, apart from 'adjective' and 'noun', pertinent numbers are far too small as to allow any relevant predictions. The other possible values of the morphosyntactic feature are 'bound form' and 'phrase' as defined in section 3.1.

Two types of simulation were run. To model the synchronic productivity profiles of the two rival suffixes, the dataset of 20th century neologisms was used both as a lexicon and as a set of test items. The parameters of the analogical program were set in such way as to ensure that, for each new word to be classified, its copy in the lexicon was excluded from the classification procedure (termed 'leave-one-out' procedure in some of the literature, cf. e.g. Daelemans & van den Bosch 2005). To model diachronic development, a dataset from an earlier century was used as a lexicon to predict the set of

neologisms for the later century. The 18th century dataset was used as a lexicon to predict the 19th century data, and the 19th century data were used as a lexicon to predict the 20th century data. The setup of the diachronic experiments was modelled on the setup used by Chapman and Skousen (2005) in their diachronic simulation study of rivalry between English negative prefixes.

The basic rationale that underlies the experiments in the present paper contains two aspects that are potentially contentious. The first is that it is pretended that the new words to be predicted by the system are true neologisms, which precludes the possibility that (simulated) speakers may already know the word. This setup was chosen to avoid circularity of the argument. If the new word is identical to an entry in the lexicon, this identical entry will automatically be the most influential member of the analogical set. Even though, depending on lexical distributions, the identical item is unlikely to be the only member of the analogical set, AM's predictions will be heavily biassed by that item. The second, more problematic aspect is that the lexicon in all simulations has been recruited from OED neologisms, with a high proportion of very low frequency words and specialised terms. It is thus clear that this type of lexicon does not constitute a realistic approximation of speakers' lexical knowledge. However, note that it is far from clear what exactly would constitute a realistic lexicon for analogy-based classification in word-formation. For example, as pointed out by a referee for this article, a conceivable alternative to the approach used in the present study would be to take the total of all *-ity* and *-ness* derivatives from all centuries as a lexicon for the 20th century simulation. The rationale that guided the choice of the lexica in the present simulation experiments was the idea that the neologisms from the same century (for the synchronic experiment) or the previous century (for the diachronic experiments) comprise a set of lexical items that can be assumed to be formally transparent members of the right morphological category (*-ity* or *-ness* derivatives, respectively) to speakers in the later century. This is not necessarily the case for items that were coined in more remote periods. Formal transparency or analysability, in turn, is a plausible key prerequisite for lexica used in analogical simulations (cf. e.g. Chapman & Skousen: 343f. for discussion). It was hoped, thus, that even if the lexicon is not an entirely realistic approximation of speakers' lexical knowledge, it could serve as a plausible proxy for that knowledge.

5.2 *Overall performance of the model*

This section will provide an overview of the overall predictive power of the simulations. For the sake of clarity, the discussion will refer to categorical AM predictions which were derived from transforming AM's probabilistic predictions into categorical values reflecting the majority vote among exemplars in the analogical set. If the predicted probability of *-ity* was 50% or higher, the predicted suffix was taken to be *-ity*. If the predicted probability of *-ity* was less than 50%, the predicted suffix was taken to be *-ness*. In addition, goodness of fit of AM's probabilistic predictions is assessed by reporting the concordance statistic *C* for each simulation (see Austin & Steyerberg 2012 for discussion). Here, the *C* statistic can be understood by taking all possible pairs of words from the lexicon used in the simulation so that in each pair, one word contains the nominaliser *-ity* and the other word contains the nominaliser *-ness. C* is the proportion of pairs in which the word containing *-ity* also has a higher predicted probability of having *-ity* than the word containing *-ness*. A *C* value of 1.0 indicates a perfect fit, and a value of 0.5 indicates a random distribution of predictions (see Baayen 2008: 204 for details).

Table 5 provides an overview of the simulation in which the 20th century neologisms were used both as the lexicon and the test set (henceforth 'synchronic simulation' for short). Predictive power is measured in terms of the percentage of correct predictions and in terms of macro-averaged F-scores. F-scores are a measure of accuracy that is frequently used in information theory (cf. e.g. Daelemans & van den Bosch 2005: 48ff. ; 78–79 for discussion). They are computed separately for each target category, in our case for *-ity* and *-ness*, as the harmonic mean of two ratios. The first ratio, termed 'recall', gives us the number of items for which the model predicts the nominaliser correctly divided by the number of items which are indeed observed to take that nominaliser. This measure is often cited as a measure of predictive accuracy. The second ratio, termed 'precision', gives us the number of items for which the model predicts the nominaliser correctly divided by the number of items for which the model predicts that nominaliser. It thus tells us how well aimed the model the model is in its predictions. Possible F-score values range between 0 and 1, with an F-score of 1 indicating faultless performance. An overall F-score is computed as the harmonic mean of the two F-scores for *-ity* and *-ness* (termed 'macro-averaged F-score').

F-score, macro-averaged:	0.88
% correct predictions (overall):	88.65%
F-score for *-ity*:	0.91
% correct *-ity*:	93.31%
F-score for *-ness*:	0.85
% correct *-ness*:	81.36%

Table 5 *The predictive power of the synchronic simulation (lexicon: 20th century, test set: 20th century, N = 564). For the probabilistic AM predictions: C = 0.92.*

As Table 5 shows, the distribution of *-ity* and *-ness* among 20th century neologisms in the OED is highly predictable, with both the overall F-score and the percentage of correct predictions clearly above 85%, and a very high concordance statistic C of 0.92. Only 64 of all 564 derivatives in the set are mispredicted by the system.

Tables 6 and 7 summarise the results of the diachronic experiments, i.e. those where the set of neologisms from the earlier century serves to predict the later century. The format of the tables is parallel to that of Table 5.

F-score, macro-averaged:	0.85
% correct predictions (overall):	85.82%
F-score for *-ity*:	0.88
% correct *-ity*:	82.56%
F-score for *-ness*:	0.83
% correct *-ness*:	90.91%

Table 6 *The predictive power of the diachronic simulation (lexicon: 19th century, test set: 20th century, $N_{test\ set}$ = 564). For the probabilistic AM predictions: C = 0.89.*

F-score, macro-averaged:	0.78
% correct predictions (overall):	79.69%
F-score for *-ity*:	0.76
% correct *-ity*:	69.85%
F-score for *-ness*:	0.80
% correct *-ness*:	86.56%

Table 7 *The predictive power of the diachronic simulation (lexicon: 18th century, test set: 19th century, $N_{test\ set}$ = 1,492). For the probabilistic AM predictions: C = 0.82.*

As can be seen especially from the averaged F-scores (0.85 for the later data, 0.78 for the earlier data) overall predictive power is very good in all diachronic experiments. Nevertheless, we see a consistent decrease in overall predictability if we go back in time (reflected also in the concordance statistic *C*, which decreases from 0.89 for the later data to 0.82 for the earlier data. This is in line with the findings presented in section 3.3, where we saw that variability between *-ity* and *-ness* decreases over time for certain types of bases.

In what follows the discussion will address in more detail the question of how exactly the system constructed the similarity relations that are relevant to this particular classification task. The focus will be on two issues: The first is how AM manages to predict the domain-specific productivity of *-ity* and the default status of *-ness* in the synchronic simulation. The second is how the increase in the domain-specificity of *-ity* that we have observed over time is reflected in the diachronic AM simulations.[13]

[13] One major aspect that is going to be ignored in the present discussion is the fact that, for adjectival bases with Germanic suffixes, *-ness* is the only option. In general, AM predictions for Germanic suffixes function like the predictions of the domain-specific productivity pattern of *-ity* to be discussed in this section. It is, however, true that for some categories AM implausibly predicts a small degree of variability between *-ness* and *-ity*. Such cases will be mentioned but not given much room in the discussion. The reason is that it is not clear whether AM's behaviour is indicative of a principled problem of the approach to predict categorical behaviour or whether it is an artifact of the experimental setup (in particular, the database and the way phonological similarity is encoded in the present simulation, cf. section 3.1 for discussion).

5.3 *Domain-specific productivity and default selection are epiphenomena*

The present section will first address the question of how AM replicates the distributions that we found in section 3.2 for the different morphological categories, and then explore the details about how productivity differences emerge in AM. Figure 9 compares the attested distribution by morphological category (in the upper panel) with AM's predictions (in the lower panel) for those morphological categories with at least 10 attestations in the data. Again, AM predictions are transformed into categorical predictions on the basis of majority votes in the analogical set. Note that, unlike in the discussion of distributions in section 3.2, Figure 9 also incorporates adjectival bases that do not contain an adjectival suffix. In contrast to accounts that base their predictions on the combinatorics of particular suffixes, AM makes predictions for all types of bases; one particularly frequent type is unsuffixed bases.

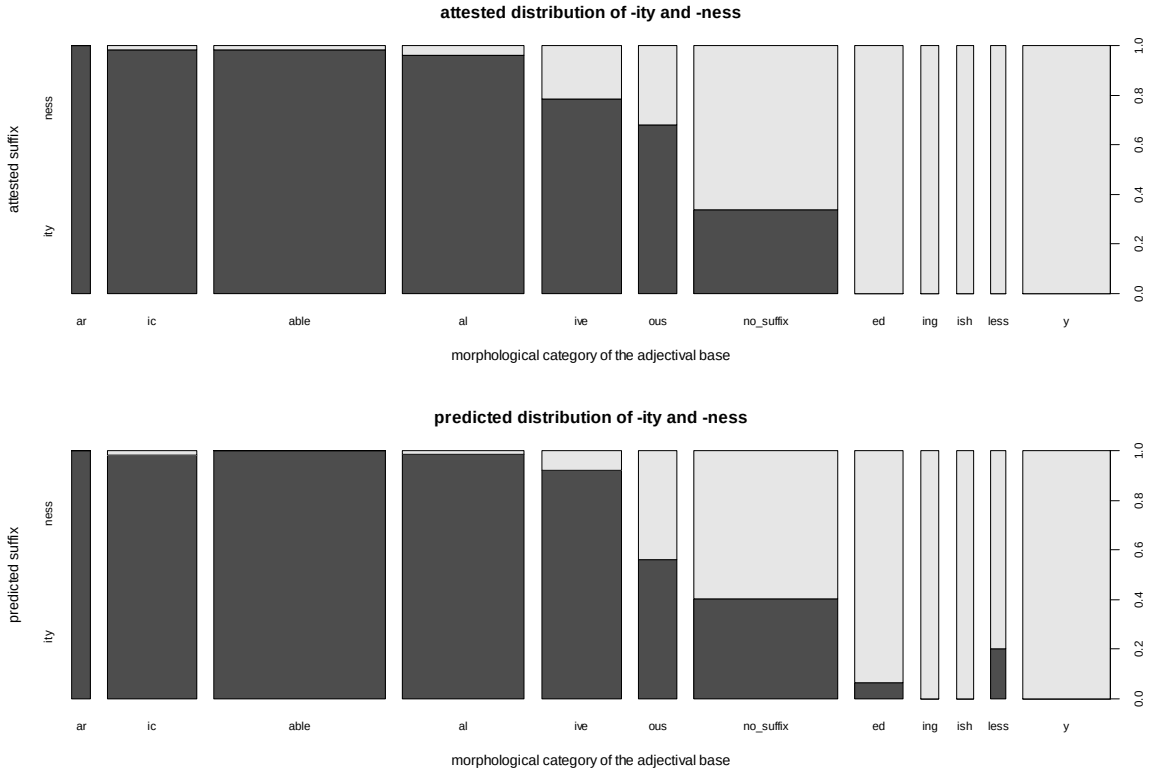

Figure 9 *Attested and predicted distributions of -ity and -ness, synchronic experiment (N = 545)*

We see that categorical AM predictions yield a predicted distribution that is highly similar to the attested distribution.[14]

As we have seen in section 3.2, two crucial phenomena characterise the productivity profiles of *-ity* and *-ness*. *-ity*'s productivity is not only concentrated in particular domains, but it also dominates these domains in comparison to *-ness*. The productivity of *-ness*, by contrast, is characterised by *-ness*'s default status. This means

[14] Cases where predictions notably differ from attested distributions concern adjectival bases ending in the Germanic suffixes *-less* and *-ed*. For these AM incorrectly predicts *-ity* for two *-ed* bases (out of 31 pertinent bases) and two *-less* bases (out of 10 pertinent bases). In all four cases the mispredictions may well be an effect of the way phonological features were coded in the present study, which is only a quite crude and incomplete approximation of phonological properties that can serve as a basis for similarity between words. The two mispredicted *-ed* bases are *oversexed* and *swollen*. In the case of *oversexed* the system had no possibility of establishing similarity with items which have a /t/ or /d/ coda of the final syllable because the coda is /kst/. In the case of *swollen* the analogical set does contain plausible analogues such as *forgottenness*, *alienness*, and *Russianness*, but these exemplars are outweighed by *-ity* bearing analogues which share with *swollen* the [əʊ] nucleus of the penultimate syllable of the base, and which happen to be quite frequent in the data set. Examples are *multipolarity*, *osmolarity*, and *molarity*. The two mispredicted *-less* bases are *prescienceless* and *jobless*. For *-less* bases the algorithm's problems in establishing similarity relations obviously lie in the fact that it has no means of distinguishing the word-final syllable /ləs/ in bases like *jobless* from that in words like *bibulous* (attested nominalisation: *bibulosity*).

that *-ness* can, in principle, occur on almost any type of base. A close analysis of the AM model reveals that this situation is in fact predicted by the algorithm. The key to understanding how this works lies in the composition of analogical sets (cf. section 4 above for an introduction of the terminology). These are significantly more local for *-ity* than they are for *-ness*. 'Locality' refers to two properties of analogical sets: One is the number of exemplars in the set, the second is the degree of similarity of members of that set with the test word. In general, these two properties are often correlated because analogical sets that consist of exemplars that are highly similar to the test word tend to be smaller than analogical sets that also incorporate less similar exemplars. In what follows the analogical sets predicting *-ity* and *-ness* will be compared in terms of both aspects of locality.

Table 8 provides analogical set sizes and measures to assess the variation in sizes for bases belonging to three morphological categories: those ending in *-able* (*-ity* leaning), *-ous* (exhibiting a relatively large degree of variation between *-ity* and *-ness*), and *-y* (*-ness* leaning). The choice of *-able*, *-ous*, and *-y* bases is arbitrary, other bases show a similar behaviour.

	-*able* bases	**-*ous* bases**	**-*y* bases**
No. of items in the morphological category:	110	25	56
analogical set size: range	2 - 127	3 - 25	6 - 31
analogical set size: median	12	9	14.5
analogical set size: mean	14.85	10.12	16.02

Table 8 *Analogical set sizes for three types of bases*

We see that for most test items, analogical sets are much smaller than the morphological category, with median or mean sizes being only a fraction of the number of pertinent items in the morphological class.[15] The second thing we see is that analogical set sizes

[15] Note, furthermore, that analogical sets are not necessarily comprised exclusively of members of the same morphological category.

are not uniform within morphological categories, exhibiting a vast range for all categories. In sum, the sizes of the analogical sets make it abundantly clear that classification in AM is not based only on the features in the phonological coding that are indicative of morphological categories, but is often much more local than that in the above-described sense.

Against this background, it is now instructive to compare *-ity* and *-ness* derivatives in terms of the degree of similarity that is constitutive of analogical sets. This will be done on the basis of the feature combinations that define the most influential gang in each analogical set that occurred in the simulation. The term 'gang' refers to a group of exemplars within an analogical set which all have exactly the same features in common (cf. especially Skousen 1995, 2002a, 2002b, 2005). These features, then, define the gang in the above-described sense, and I will henceforth refer to them as 'constitutive features' of the gang. The 'most influential gang' is the gang that has the strongest weight and, thus, the strongest influence on the classification of a new item. Figure 10 provides an illustration of the gang terminology for the new word *aggressivity* with its analogical set. The general format is identical to that of Figure 8 above, the only difference is that gangs are shaded, and that gang members are represented only in terms of the constitutive features of the gang.

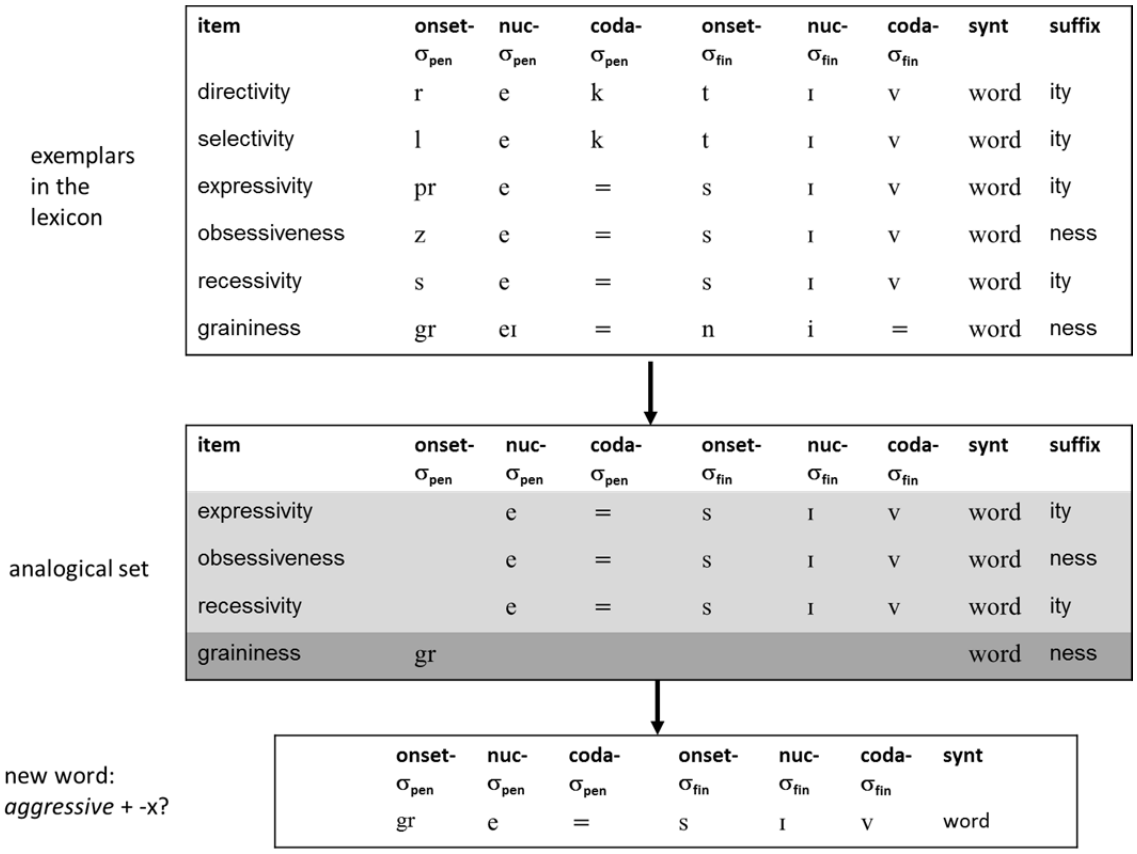

Figure 10 *Gangs in the analogical set*

The analogical set comprises four exemplars, *expressivity*, *obsessiveness*, *recessivity*, and *graininess*, which are split into two gangs. The first gang consists of *expressivity*, *obsessiveness*, and *recessivity*. The constitutive features of that gang are the phonological features nuc-σ_{pen}, coda-σ_{pen}, onset-σ_{fin}, nuc-σ_{fin}, and coda-σ_{fin}, as well as the syntactic category. The second gang in Figure 8 consists of only one exemplar, *graininess*. The base *grainy* is only vaguely similar to *aggressive*, the constitutive features of the gang being onset-σ_{pen} and the syntactic category. In terms of their influence on classification of *aggressive*, the algorithm gives much more weight to the first gang than to the second gang. This is due to two reasons: One is that the first gang is defined by more constitutive features, that is, exemplars in the gang are more similar to the new word than the exemplar in the second gang. The second is that the first gang has more members than the second gang.[16]

[16] Technically, this is achieved because AM computes weights for each exemplar based on the number of disagreements between the features of the exemplar and the features of the new word for each possible context, where 'context' is defined as any possible combination of features. More similar exemplars are given more weight because they appear in more contexts than less similar exemplars. The details of

The question now is whether there is a principled difference between *-ity* and *-ness* words in terms of the similarity structure in their analogical sets. Figure 11 provides an overview of the number of constitutive features that define the most influential gang for each *-ity* and *-ness* derivative in the test set.

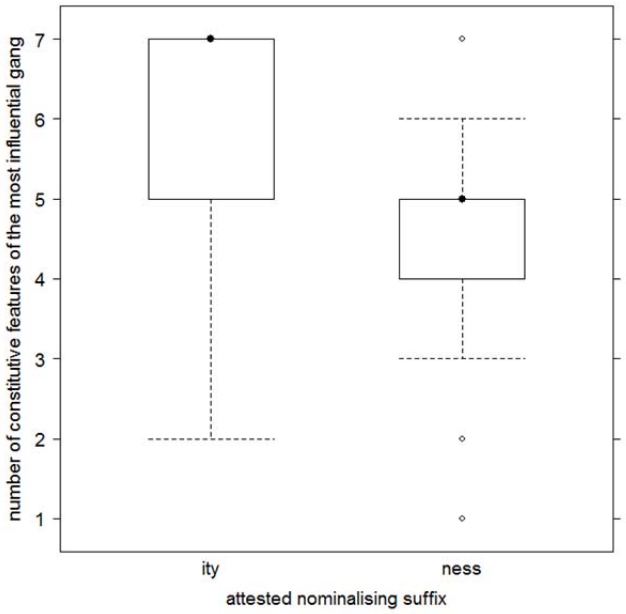

Figure 11 *Similarity between the most important gang and the test word, N = 564*

The boxplot shows that there are clear differences between the similarity structure of analogical sets of *-ity* and *-ness* derivatives. The median number of features shared by the most influential gang with the test word is 7 for *-ity* derivatives and 5 for *-ness* derivatives (indicated by the black dots in the plot). The boxes indicate that for 75% of all *-ity* derivatives, the range of the number of shared features is between 5 and 7, whereas for 75% of all *-ness* derivatives it is between 4 and 5. As is clear from the location of the mean with respect to the interquartile range represented by the boxes, the pertinent means are somewhat lower than the medians for both groups of derivatives (for *-ity* derivatives: 6.17; for *-ness* derivatives: 4.70). A Wilcoxon test shows that the difference between the means is indeed highly significant ($W = 61003$, $p < 2.2e\text{-}16$).[17]

weight computation, which do not need to concern us here, are discussed, for example, in Skousen (2002a), Skousen (2002b), and Skousen (2005). In the example in Figure 10, the first gang is given 90% of the overall vote on nominaliser selection, and the second gang is given 10%.

[17] Note that the observed difference is not an artifact of the fact that two *-ity* leaning basal suffixes are disyllabic. These are *-able* and *-ary*. The effect stays robust if we remove all *-able* and *-ary* items from the

Tables 9 and 10 provide representative examples of the similarity structure of analogical sets for *-ity* and *-ness* derivatives (for the new words *aromaticity* and *mopiness*, respectively).

constitutive features of the gang							**exemplars in the gang**	**weight of gang vote**
Ons_{pen}	**Nuc_{pen}**	**Cod_{pen}**	**Ons_{fin}**	**Nuc_{fin}**	**Cod_{fin}**	**Synt**		
m	æ	=	t	ɪ	k	word	*monochromaticity* *systematicity*	67.3%
	æ			ɪ	k	word	*aeroelasticity* *monosyllabicity* *oceanicity* *syllabicity*	15.0%
	æ		t	ɪ	k	word	*facticity* *tacticity*	11.2%
		=	t	ɪ	k	word	*analyticity*	3.7%
m	æ			ɪ		word	*massivity*	2.8%

Table 9 *Gangs in the analogical set for aromatic (⬧ aromaticity)*

dataset, even though there is a slight decline in the mean of features constitutive of *-ity* gangs (W = 37685.5, p < 2.2e-16, N = 449).

constitutive features of the gang							exemplars in the gang	weight of gang vote in the overall vote
Ons_{pen}	**Nuc_{pen}**	**Cod_{pen}**	**Ons_{fin}**	**Nuc_{fin}**	**Cod_{fin}**	**Synt**		
		=	p	i	=	word	*cuppiness, loppiness, nippiness, peppiness, snipiness, stripiness*	32.4%
	əʊ	=	p	i	=	word	*dopiness*	16.2%
m		=		i	=	word	*mateyness, mininess*	14.4%
m	əʊ	=				word	*multimodality neomodality*	9.0%
	əʊ	=		i	=	word	*nosiness phoniness*	7.2%
m				i	=	word	*maltiness supermanliness*	7.2%
m	əʊ	=			=	word	*molarity osmolarity*	7.2%
		=	p		=	word	*peaness moneyness*	3.6%
			p	i	=	word	*frumpiness*	2.7%

Table 10 *Gangs in the analogical set for mopey (⬧ mopiness)*

An important structural difference between the analogical sets for the bases *aromatic* and *mopey* is how similar the most influential exemplars in the analogical set are to the new word. The most influential gang for *aromatic* shares all seven coded features with the new word. As we saw in Figure 11, the same is true for half of the *-ity* derivatives in the simulation. Even though the gang consists of only two exemplars, it accounts for

67.3% of the overall vote on the predicted nominaliser for *aromatic*. The situation is different for *mopey*, for which *-ness* is predicted as a nominaliser. Here the most influential gang shares only four features with the new word.[18] Overall, we also see that the analogical set for *mopey* contains more gangs than the set for *aromatic*, where gangs differ more widely in terms of their constitutive features.

The findings of the analysis thus show that differences in productivity profiles between *-ity* and *-ness* emerge from differences in the similarity structure that is relevant for the classification in the lexicon. For *-ity* derivatives, classification is very local, i.e. dominated by highly similar exemplars. The effect is the niche behaviour that was observed for *-ity* in section 3 of this paper as well as in related research (esp. in Lindsay 2012). For *-ness* derivatives, by contrast, classification is relatively more general than for *-ity*, i.e. dominated by less similar exemplars. The effect is that relevant similarity relations tend to be more sparse, making the behaviour of *-ness* look like default selection, i.e. a kind of selection that is not primarily dependent on specific feature constellations.

These results are fully in line with recent work on productivity in computational analogical frameworks (esp. Albright & Hayes 2003, Keuleers et al. 2007, Keuleers 2008). The classic phenomenon that has attracted most interest here is English past tense formation, specifically, the insight that, contrary to the expectations of many grammatical models, not only regular, but also allegedly irregular formation patterns are productive, albeit to different degrees (Bybee & Moder 1983, Prasada & Pinker 1993). This has given rise to a vivid theoretical debate about the status of regular and irregular processes in morphology especially between dual-mechanism models (Pinker 1991, Marcus et al. 1992, Prasada & Pinker 1993, Pinker & Prince 1994), connectionist models (McClelland & Rumelhart 1985, Matthews 2013), and analogical models (Derwing & Skousen 1994, Eddington 2000, Keuleers 2008, Chandler 2010). What is interesting in the context of the present investigation is that the differences between irregular and regular past tense formation in terms of their productivity profiles are in principle similar to those between *-ity* and *-ness*. Like *-ity*, irregular past tense formation

[18] The reason why this gang is given more weight than the second most important gang in Table 10 (the exemplar *dopiness*), which shares five features with the new word, is that it has many more members and by these means compensates for the greater weight that *dopiness* has by virtue of its greater similarity with the new word.

exhibits niche productivity in the sense that its productivity is mainly confined to bases that are highly similar to existing irregular bases. Like *-ness*, regular past tense formation has a status that looks like a default. For the past tense this situation, which has been a classic source of evidence for dual mechanism models, has been shown to also emerge in a single-mechanism analogical model (Derwing & Skousen 1994, Eddington 2000, Keuleers 2008, Chandler 2010). The way in which it emerges is stunningly similar to what we saw for *-ity* and *-ness*: Domain-specific behaviour ('niches' in the terminology of Lindsay 2012, Lindsay & Aronoff to appear) in an analogical model translates into strong gang behaviour exerted by lexical items that are highly similar to a test item. 'Default' translates into influence from (relatively) more distant items that are also usually more sparsely distributed over the similarity space. The metaphors 'niche' and 'default' constitute endpoints on a gradient scale.

5.4 *Predicting change: the role of certainty and strong gangs*

This section will compare the two diachronic simulation experiments with the synchronic model. The question to be addressed is how exactly the models represent the fact we observed in section 3.3 that for bases with frequent Latinate suffixes, *-ity* becomes the dominant nominaliser whereas *-ness* retreats from those domains. Again the morphological category of the base will be used as a convenient label for those domains, without much theoretical commitment. As we saw in section 5.3 for the synchronic study, analogical sets in the present simulations are only broadly correlated with morphological categories.

The situation that *-ity* comes to robustly dominate contexts where before there was variation between *-ity* and *-ness* can be seen in two properties of the AM models: One is that, as we progress in time, AM predicts a decreasing amount of variability between the two nominalisers for these contexts. The second is that in time we see the emergence of strong gangs favouring *-ity*. These gangs are strong in the sense that they reappear in several analogical sets. Both properties will be discussed in turn.

In order to look at the predicted variability in pertinent domains (i.e. bases with frequent Latinate suffixes), we need to consider AM's probabilistic predictions for nominalisers rather than the categorical predictions that have been in the focus far. Recall from section 4 that AM predictions are inherently probabilistic. To avoid unnecessary complexities, these were in the discussion so far transformed into

categorical predictions reflecting majority choices within analogical sets. The three boxplots in Figure 12 plot the predicted probabilities of *-ity* for pertinent bases with Latinate suffixes. The upper panel represents the simulation in which the 19th century data were predicted on the basis of the 18th century lexicon. The mid panel represents the simulation in which the 20th century data were predicted on the basis of the 19th century lexicon. The lower panel contains the synchronic experiment for comparison, in which the 20th century data were predicted on the basis of the 20th century lexicon.

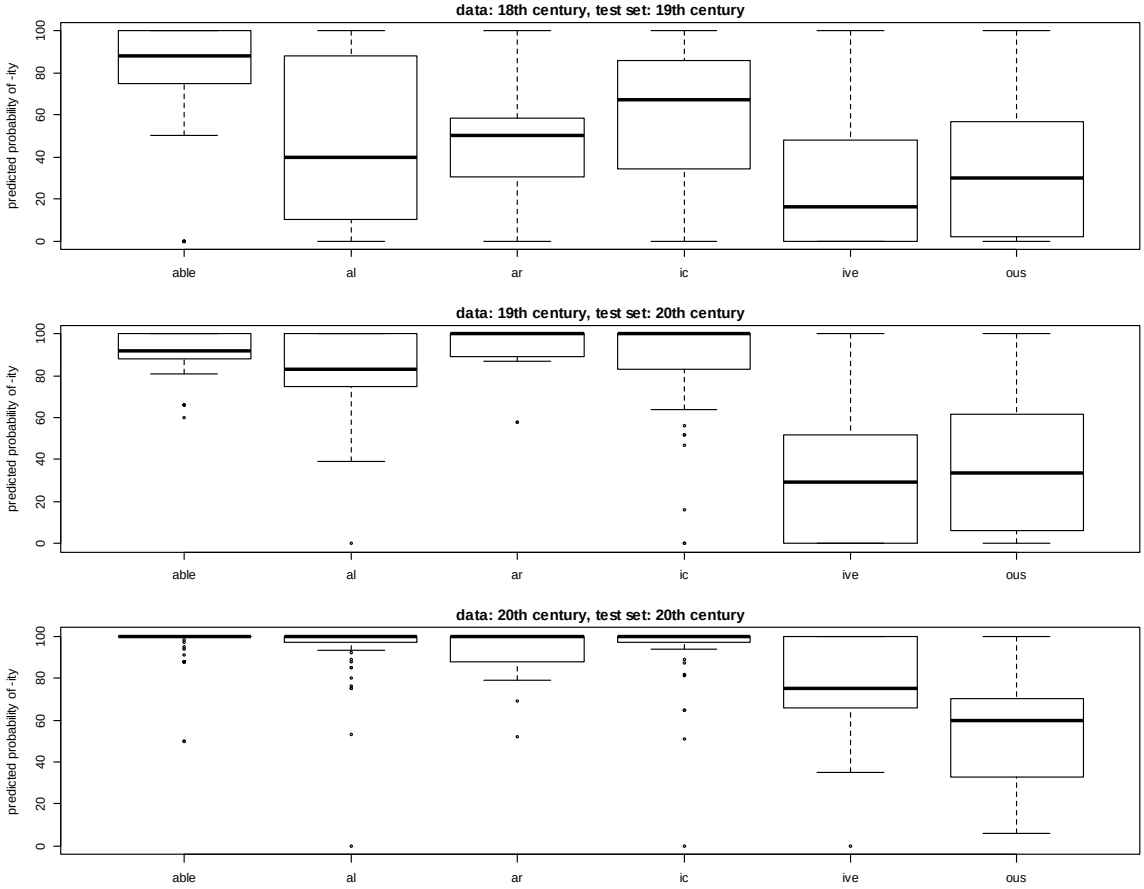

Figure 12 *Predicted probability of -ity for bases with frequent Latinate suffixes in the diachronic experiments and in the synchronic experiment*

The boxplots show two things. First of all, for all morphological categories, the median predicted probability of *-ity*, which is represented as the line in each of the boxes, increases over time. This median closely corresponds to the proportion of *-ity* derivatives in the pertinent data set that is used as a lexicon in the simulation (cf. section 3.3 for the pertinent distributions). Secondly, the boxplots show that the range of

predicted probabilities becomes smaller over time. This is reflected in the fact that the boxes, which represent 75% of the data, become smaller. What this means is that *-ity* predictions become more and more robust, and *-ness* becomes less and less likely for those morphological categories.

In order to see how gang behaviour produces the observed patterns, we need to analyse the simulations in terms of which gangs are most active in classification. To this end, we take a look again at the most influential gangs in the analogical sets of *-ity* derivatives (cf. section 5.3 for an explanation of the terminology). The simulation that predicts the 19th century data on the basis of the 18th century data will be used as an illustrative example. Figure 13 is a so-called Fruchterman-Reingold diagram (Fruchterman & Reingold 1991), as implemented in the R package igraph (Csardi & Nepusz 2006). Fruchterman-Reingold diagrams are commonly employed for social network analysis. The purpose of such a graph is to visualise relationships between members of a network in an undirected graph. The graph represents these relations by means of analogy with natural forces. Thus, related members are subject to attractive forces, whereas unrelated members are subject to repulsive forces. In our case, we will consider the *-ity* derivatives in the 19th century test set to be the members of the network. Relationships between the members of the network are defined in terms of the set of features that are constitutive of the most influential gang in the simulation that predicts the nominaliser for those derivatives (in this case, on the basis of the 18th century data). The effect is that derivatives which are classified on the basis of the same most influential gang will cluster together in the graph. Notice that the clustering visible in the graph is actually less strong than in reality, given that the graph can only take into account the most influential gang. If all gangs in the analogical set would be taken into account, overlap between gangs would be even greater.

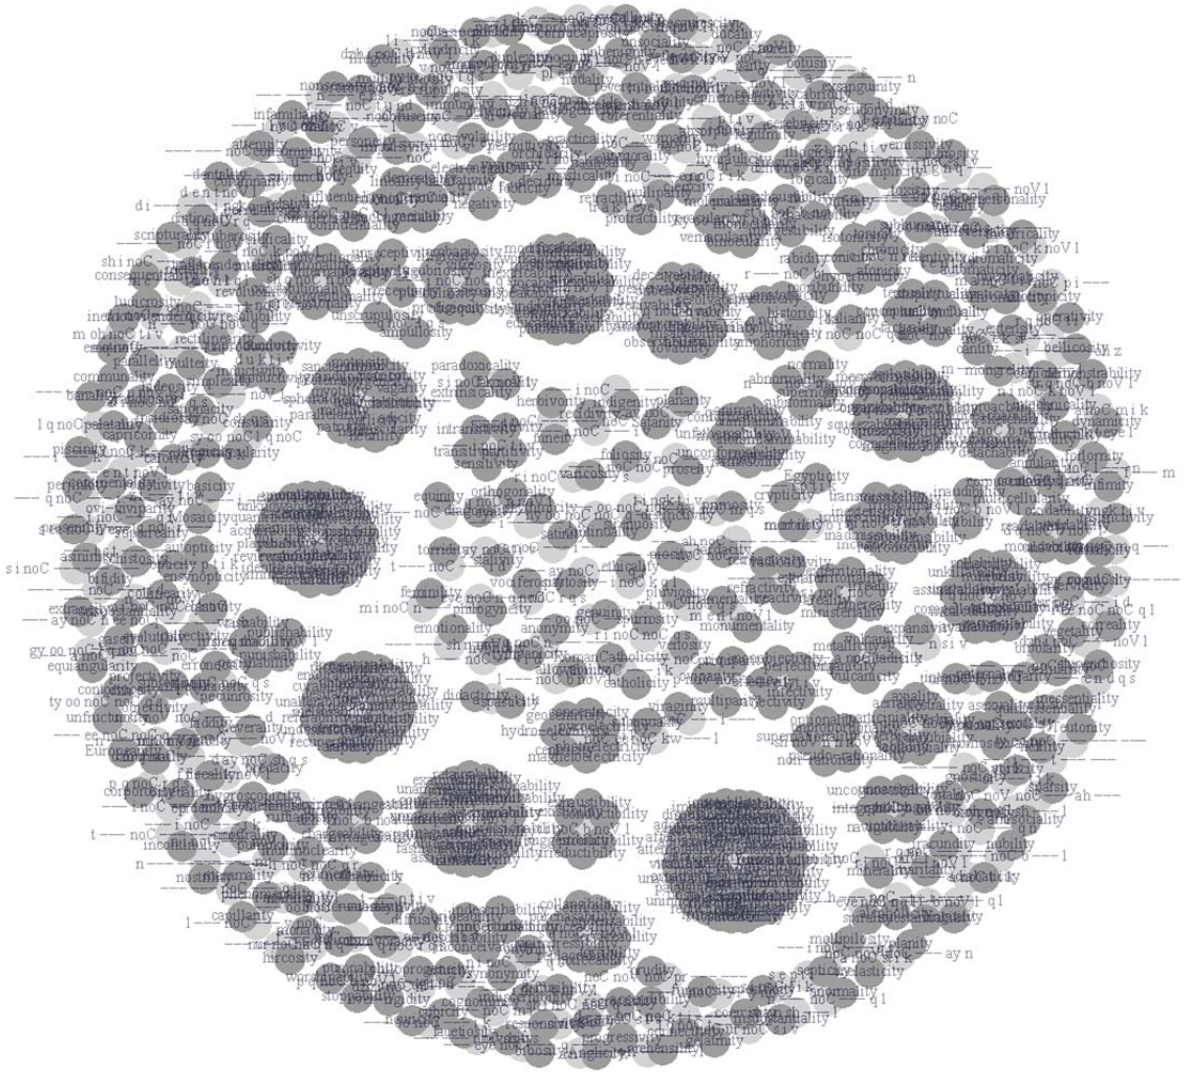

Figure 13 *Derivatives and the constitutive features of their most influential gangs for -ity derivatives; data set: 18th century, test set: 19th century*

Each dark grey node in Figure 13 represents an *-ity* derivative in the 19th century test set. Each light grey node represents a gang that acts as the most influential gang for some *-ity* derivative. There are two possible configurations: If the most influential gang for a derivative is only active in that derivative, there is a pair of one dark grey and one light grey node in the graph. If, however, a gang is active for more derivatives, there will be a cluster of dark grey nodes around the relevant light grey gang node. The most striking property of the graph is that most dark nodes are clustered, and there are several very large clusters. If we look at the largest clusters, which may be defined here as clusters comprising at least seven dark grey nodes, we see that 12 of the 15 clusters of

that size comprise *-able* derivatives. The most influential gang of such derivatives typically consists of highly similar exemplars, which share all seven coded features. An example is the biggest cluster in the lower right part of the graph, which comprises 53 derivatives whose bases end in [təbl̩] and whose syntactic category is 'word'. Examples of pertinent derivatives are *affectability*, *adoptability*, *attemptability*, and *reductibility*. The three big clusters that are not related to *-able* comprise gangs

- where the base-final syllable ends in [ɪk] and the nucleus of the penultimate syllable is [e] (as in *concentricity*, *geocentricity*, *hydroelectricity*, and *photoelectricity*),
- where the base is bound (as in *biplicity*, *caducity*, and *tenebrity*), and
- where the base ends in [ɪəs] (as in *vitreosity*, *prodigiosity*, *lugubriosity*, and *ambagiosity*).

The situation that the same gang is active in many derivatives, i.e. the situation that is represented by the clusters in the graph, is the prerequisite both for patterns to emerge and for change to happen in nominaliser selection. For example, as we have seen in section 3.3, the 19th century saw a sharp rise in nominalisations based on *-able* derivatives. Within that set, for example, the number of bases ending in the sequence [nəbl̩] rose from 4 in the 18th century data (*alienability*, *inalienability*, *ponibility*, *interponibility*) to 29 in the 19th century (e.g. *retainability*, *unamenability*, *assignability*, *fashionability*). Since all four 18th century [nəbl̩] bases take *-ity* as a nominaliser, these four exemplars, acting as a gang for 29 new words, exerted a disproportionately strong pressure towards *-ity* among *-able* bases. The 29 derivatives taking [nəbl̩] as the most influential gang can be seen as a big cluster on the bottom left side of the circle of big clusters in Figure 13.

As another example of the role of strong gangs in promoting change, we may consider the behaviour of adjectival bases ending in *-al* in the earliest simulation (dataset: 18th century, test set: 19th century). For *-al* bases, we have seen in Figure 12 above that there is large variation in terms of the probability of *-ity* predicted by AM. This is captured in the box-and-whiskers plot in Figure 12 by the dimensions of the box for *-al* suffixed bases, which shows that for 75% of *-al* suffixed bases, the predicted probabilities of *-ity* are between 9.25% and 87.5%. The reason is that for AM *-al* derivatives do not form a homogeneous class apart from formal similarities in the base-final syllable. Given that analogies for the Latinate bases are very local,

predictions vary with the exact phonemic makeup of the two base-final syllables. For example, AM makes radically different predictions for bases ending in /tjʊəl/ and /rɪəl/. The 18th century lexicon contains only one base ending in /tjʊəl/: *eventual* (♦ *eventuality*). This base serves as a strong gang with only one member favouring *-ity* for all 19th century test data ending in /tjʊəl/. The 19th century test set contains 6 pertinent bases, which appear as a cluster in the upper right corner of Figure 13: *conceptuality*, *factuality*, *textuality*, *unpunctuality*, and *accentuality*. By contrast, analogical sets for bases ending in /rɪəl/ provide stronger support for *-ness* than for *-ity*. The 18th century lexicon contains five pertinent words: *corporealness*, *etherealness*, *imperialness*, *territoriality*. For the ten test items from the 19th century test set which end in /rɪəl/, then, the predicted probability of *-ity* is relatively small compared to *-ness*. The pertinent items are *dictatorial* (*-ness*), *mercurial* (*-ness*), *ministerial* (*-ness*), *aerial* (*-ity*), *areal* (*-ity*), *ethereal* (*-ity*), *exterritorial* (*-ity*), *extraterritorial* (*-ity*), *irreal* (*-ity*), and *ministerial* (*-ity*).

5.5 *Section summary*

This section reported on a series of simulation experiments that explores how AM predicts synchronic variation and diachronic change in *-ity* / *-ness* selection on the basis of merely the phonological makeup of the two base-final syllables and a rudimentary specification of word-class information. In terms of overall predictive power, it was shown that this model accounts accurately for between about 88% and 78% of the data. Highest F-scores were reached for the 20th century data, and predictive power shows a consistent decline in the simulations that are concerned with the two earlier centuries. The success of the model suggests that formal factors like the ones coded indeed play a paramount role in the rivalry between the two nominalisers, and that their distribution is highly predictable on the basis of these factors already. Furthermore, the comparison of predictive power reached in the diachronic simulations shows that predictability of *-ity* and *-ness* selection is increasing over time. This is reflected both in the overall accuracy scores of the model based on majority choices in the analogical set and the ranges that can be seen in the probability distributions within analogical sets.

This finding lends strong support to the idea that choices between nominalisers are determined online on the basis of lexical distributions, and that probabilistic biasses in those distributions can lead to a slow decrease in variability over time. On the

surface, then, this situation looks like a type of regularisation of the system. However, 'regularisation' does crucially not mean here that there is an underlying rule, but is only used here to refer to the emergence of clear preference patterns.

A detailed analysis of the synchronic simulation experiment with the 20th century data revealed how differences in the productivity profiles of *-ity* and *-ness* emerge from the model. Crucially, it was shown that differences in the productivity behaviour of the two nominalisers are an effect of differences in the locality of classification in the model. Classification is more local in AM than in other approaches which are based on the morphological category of the base. Furthermore, *-ity* predictions are significantly more local than *-ness* predictions. This difference between the two nominalisers emerges automatically in the AM model because it can be read directly from the lexical distribution: exemplars taking *-ity* are, overall, more similar to each other than exemplars taking *-ness*.

Furthermore, two diachronic simulation studies were discussed which revealed how the change in preference patterns for *-ity* is represented in the analogical model. It was argued that the prediction of change rests on two crucial properties of the model: One is the probabilistic nature of predictions, the second one is the locality of prediction. The probabilistic nature of predictions makes sure that variability persists, but is slowly weakened in time. In the simulation studies this was reflected by the decrease in the range of predicted probabilities for the later data in comparison to the earlier data. The effect of locality is that change can come about if the distribution of pertinent properties in the lexicon changes, enhancing the activity of strong gangs while weakening others.

6 CONCLUSION

The present paper has sought to make a contribution to our understanding of the competition between *-ity* and *-ness* both on an empirical and a theoretical level. The quantitative survey of the distribution of the two suffixes among OED neologisms from the three past centuries has helped to shed light on two issues that are at the centre of the debate in the current literature on the two suffixes. One is their productivity profiles, the other is the diachronic emergence of these profiles. With respect to the former, the analysis has shown that the distribution of the two suffixes is semi-complementary, in

the sense that, as is well known, the two suffixes are subject to different selectional restrictions for some types of bases, but may select bases from the same catgory for other types of bases. Examples of complementary restrictions are those ruling out *-ity* for bases with Germanic adjectival suffixes or those ruling out *-ness* for bound bases. More interestingly, however, we saw that among bases which in principle allow both *-ity* and *-ness*, there are clear preference patterns, most of which clearly favour *-ity* over *-ness*. These are bases ending in the Latinate adjectival suffixes *-able*, *-ar*, *-al*, *-ic*, and *-ive*. With the exception of *-ive* bases, *-ity* preference is near-categorical for such bases. The validity of the findings gathered from the OED data is substantiated by the fact that they are fully in line with observations found in earlier corpus-based work (esp. Baayen & Renouf 1996, Lindsay 2012).

This means, then, that differences between *-ity* and *-ness* selection cannot be accounted for in terms of appeals to global differences in productivity between the two suffixes, of the type often postulated especially in stratum-based theories (e.g. Kiparsky 1982a, Kiparsky 1982b, Giegerich 1999). The reason is that such models do not provide a straightforward way of accounting for the interdependence of the productivity profiles of the two suffixes. Instead, differences are more adequately captured in terms of the scope of domains in which the two suffixes are productive. Such scope is more particularised for *-ity* than for *-ness*, a situation which has been described as niche productivity in the case of *-ity* (Lindsay 2012) and a default behaviour in the case of *-ness* (Bauer et al. 2013). Crucially, domains are interdependent. A natural explanation for such interdependence is provided by the idea that speakers' decisions about which of the two suffixes they choose are based on the distribution of similar items in the lexicon, i.e. on analogical inferencing.

In terms of the diachronic emergence of productivity patterns, the analysis in the present study was the first to investigate the recent past of the development in Late Modern English, a period that has not been in the focus of much work in historical word-formation in general (cf. Dietz in press). It was shown that the productivity profiles that we see in the contemporary data are the result of an ongoing and consistent development throughout that period, which is characterised by a constant increase of the productivity of *-ity* in the domains in which it occurs. This is accompanied by a corresponding decrease of the productivity of *-ness* in those domains. The diachronic

facts therefore provide further evidence that competition between *-ity* and *-ness* in language use involves analogical reasoning. If the lexical distribution influences speakers' online choices, such choices are necessarily probabilistic. Choices based on these probabilities are bound to develop towards what looks like a regularisation of the system. The effect is the apparent stabilisation of *-ity* in the domains in which it occurs. Lindsay has convincingly demonstrated how such a development can be triggered by lexical change accompanying language contact in Middle English (Lindsay 2012). The present investigation has shown that the development is much more persistent than hitherto assumed, which implies that analogical mechanisms are not only relevant in situations of radical changes in the system (such as those in Middle English). The exact relation between the observed developments and lexical change in Late Modern English needs to be explored in future research.

The simulation studies with AM showed how an analogical model can predict preference patterns for *-ity* and *-ness* solely on the basis of the phonological makeup of the two base-final syllables and rudimentary information about the syntactic category of the base. We saw that AM successfully predicts the attested distribution in all experiments. A close analysis of the model has revealed how the key insights gained from the analysis of the OED data are accounted for by the analogical model. From this model, testable predictions emerge about the nature of productivity, variation and change.

With respect to productivity, we saw that differences in the degree of specificity of domains for *-ity* and *-ness* translate into differences between local and more general analogies. Whereas classification in AM is quite local in general, the default status of *-ness* as compared to *-ity* emerges as an epiphenomenon of the fact that analogies predicting *-ness* are based on a lesser degree of similarity than those predicting *-ity*. Crucially, this effect emerges because AM uses information about the distribution of *-ity* and *-ness* in the lexicon to determine analogical set composition (via 'homogeneity', cf. Skousen 2002a, 2002b, 2005 for detailed explanation). With respect to change we saw that AM's prediction is that change is connected to the probabilistic and to the local nature of classification. Specifically, change can occur if the distribution of relevant properties in the new word stock is different from that in the older lexicon, leading to a strengthening of the activity of relevant gangs and a weakening of others.

This was illustrated with examples from the simulation studies conducted in the present paper. However, it is a task for future research to test this prediction in detail. This was not possible within the limits of the present study. Such test must crucially involve a more realistic representation of lexical distributions both in terms of the coding of features relevant to describe formal similarity between bases and in terms of the question of the composition of the lexicon in general.

In terms of the debate in the literature about the role of morphological vs. purely phonological information for affix selection (discussed in section 2.1), the simulations in the present paper not only show that very good predictive power can be achieved without including the morphological category of the base as a predictor. The analysis of the AM models has also demonstrated that AM's success crucially relies on the model having access to more specific and finegrained information about the phonological form of the base than there is in the morphological category alone. The reason is that this more specific information allows AM to refer to more finegrained differences in similarity between exemplars in the lexicon. As we have seen, it is these finegrained differences that allow the key properties of an analogy-based model of productivity, variation, and change to emerge: analogies of different degrees of locality (cf. section 5.3) and the emergence of strong gangs (cf. section 5.4). It is also clear, however, that this does not necessarily mean that morphological information is irrelevant. It may be that adding morphological information on top of the features already coded would enable AM to make even more finegrained distinctions in computing similarity relations. The same holds for the inclusion of other types of information, such as semantic or pragmatic information. Future research may show how much can be gained from adding further information. Crucially, this must be done with a different dataset as in the present model there is not much room for predictive improvement to make a convincing case here

7 APPENDIX

morphological class	**N**	**%**	**examples**
no suffix	160	24.92%	queerity, complexness
-able	124	19.31%	disagreeability, approachableness
-al	78	12.15%	squirality, chronicalness
-ous	76	11.84%	nebulosity, amphibiousness
-y	52	8.10%	soapiness
-ive	36	5.61%	cogitativity, oppressiveness
-ent	23	3.58%	prolificity, causticness
-ic	22	3.43%	emergentness
-ing	19	2.96%	shockingness
-ed	16	2.49%	detachedness
-ish	14	2.18%	oysterishness
-ical	12	1.87%	analogicalness
-ar	10	1.56%	globularity, angularness
Total	642	100.00%	

Table A1 *Morphological classes of bases, 18th century*

morphological class	**N**	**%**	**examples**
-able	313	21.47%	adoptability, bearableness
no suffix	271	18.59%	genuinity, defunctness
-y	203	13.92%	jokiness
-ive	152	10.43%	adaptivity, accumulativeness
-al	130	8.92%	areality, centralness
-ous	89	6.10%	erroneosity, incautiousness
-ic	80	5.49%	caloricity, chaoticness
-ed	67	4.60%	causedness
-ing	33	2.26%	annoyingness
-less	32	2.19%	beardlessness
-ly	26	1.78%	churchliness
-ish	26	1.78%	moodishness
-ar	21	1.44%	planarity
-ful	15	1.03%	fitfulness
Total	1458	100.00%	

Table A2 *Morphological classes of bases, 19th century*

syntactic category	***-ity* derivatives**		***-ness* derivatives**	
	N	**%**	**N**	**%**
adjective	655	89.4%	713	94.0%
noun	23	3.1%	13	1.7%
bound form	51	7.0%	0	
phrase	0		5	0.7%
compound	1	0.1%	13	1.7%
other categories (adverb, verb, preposition, pronoun, suffix)	3	0.4%	15	2.0%
Total	733	100%	759	100%

Table A3 *-ity / -ness derivatives by syntactic category of the base, 19th century*

syntactic category	***-ity* derivatives**		***-ness* derivatives**	
	N	**%**	**N**	**%**
adjective	275	89.9%	397	97.3%
noun	9	2.9%	7	1.7%
bound form	18	5.9%	0	
phrase	1	0.3%	1	0.2%
compound	0		1	0.2%
other categories (adverb, verb, preposition, pronoun, suffix)	3	1.0%	2	0.4%
Total	306	100%	408	100%

Table A4 *-ity / -ness derivatives by syntactic category of the base, 18th century*

8 REFERENCES

Albright, Adam. 2009. Modeling analogy as probabilistic grammar. In James P. Blevins & Juliette Blevins (eds.), *Analogy in Grammar*, 185–213. Oxford: OUP.

Albright, Adam & Bruce Hayes. 2003. Rules vs. Analogy in English Past Tenses: A Computational/Experimental Study. *Cognition* 90. 119–161.

Anshen, Frank & Mark Aronoff. 1981. Morphological productivity and phonological transparency. *Canadian Journal of Linguistics* 26. 63–72.

Anshen, Frank & Mark Aronoff. 1988. Producing morphologically complex words. *Linguistics* 26. 641–655.

Anshen, Frank & Mark Aronoff. 1999. Using dictionaries to study the mental lexicon. *Brain and Language* 68. 16–26.

Anttila, Raimo. 2003. Analogy: the warp and woof of cognition. In Brian D. Joseph & Richard D. Janda (eds.), *The Handbook of Historical Linguistics*, 425–440. Malden, MA: Blackwell.

Arndt-Lappe, Sabine. in press. 47. Word-formation and analogy. In Peter O. Müller, Ingeborg Ohnheiser, Susan Olsen & Franz Rainer (eds.), *Word-Formation - An International Handbook of the Languages of Europe*. Berlin: de Gruyter Mouton.

Arndt-Lappe, Sabine & Melanie Bell. 2013. Analogy and the nature of linguistic generalisation: locality, generality, and variability in English compound stress. submitted for publication.

Aronoff, Mark. 1976. *Word Formation in Generative Grammar*. Cambridge, MA: MIT Press.

Austin, Peter C. & Ewout W. Steyerberg (2012). Interpreting the concordance statistic of a logistic regression model: relation to the variance and odds ratio of a continuous explanatory variable. *BMC Medical Research Methodology* 12: 82.

Baayen, Harald R. 2008. *Analyzing Linguistic Data*. Cambridge: CUP.

Baayen, Harald R. & Anneke Neijt (1997). Productivity in context: a case study of a Dutch suffix. *Linguistics* 35 (3): 565–587.

Baayen, Harald R. & Antoinette Renouf. 1996. Chronicling the Times: Productive lexical innovations in an English newspaper. *Language* 72(1). 69–96.

Baeskow, Heike. 2012. -Ness and -ity: Phonological exponents of n or meaningful nominalizers of different adjectival domains? *Journal of English Linguistics* 40(1). 6–40.

Bauer, Laurie. 2001. *Morphological Productivity*. Cambridge: CUP.

Bauer, Laurie, Rochelle Lieber & Ingo Plag. 2013. *English Morphology: A Reference Guide to Contemporary English Word-Formation and Inflection*. Oxford: OUP.

Beal, Joan C. 2004. *English in modern times: 1700-1945*. London: Routledge.

Beal, Joan C. 2012. 5 Periods: Late Modern English. In Alexander Bergs & Laurel J. Brinton (eds.), *English Historical Linguistics*, 63–78. Berlin: Mouton de Gruyter.

Becker, Thomas. 1990. *Analogie und morphologische Theorie* (Studien zur theoretischen Linguistik). München: Fink.

Bermúdez-Otero, Ricardo. 2012. The Architecture of Grammar and the Division of Labour in Exponence. In Jochen Trommer (ed.), *The Phonology and Morphology of Exponence - the State of the Art*: *to appear.*. Oxford: OUP.

Bermúdez-Otero, Ricardo & April M. McMahon. 2006. English phonology and morphology. In Bas Aarts & April M. McMahon (eds.), *The Handbook of English Linguistics*, 382–410. Oxford: Blackwell.

Booij, Geert. 2010. *Construction Morphology*. Oxford: Oxford University Press.

Bybee, Joan & Carol L. Moder. 1983. Morphological classes as natural categories. *Language*(59). 251–270.
Chandler, Steve. 2010. The English past tense: Analogy redux. *Cognitive Linguistics* 21(3). 371–417.
Chapman, Don & Royal Skousen. 2005. Analogical modeling and morphological change: the case of the adjectival negative prefix in English. *English Language and Linguistics* 9(2). 333–357.
Chomsky, Noam & Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.
Csardi, Gábor & Tamás Nepusz. 2006. The igraph software package for complex network research. *InterJournal* Complex Systems. 1695.
Daelemans, Walter & Antal van den Bosch. 2005. *Memory-Based Language Processing*. Cambridge: CUP.
Daelemans, Walter, Jakub Zavrel, Ko van der Sloot & Antal van den Bosch. 1999 et seq. *TiMBL: Tilburg Memory Based Learner* : Available from http://ilk.uvt.nl/timbl/.
Derwing, Bruce I. & Royal Skousen. 1989. Morphology in the mental lexicon: a new look at analogy. In Geert Booij & Jaap van Marle (eds.), *Yearbook of Morphology 1989*, 55–71. Dordrecht: Foris.
Derwing, Bruce I. & Royal Skousen. 1994. Productivity and the English past tense: Testing Skousen's analogical model. In Susan D. Lima, Roberta Corrigan & Gregory K. Iverson (eds.), *The Reality of Linguistic Rules*, 193–218. Amsterdam / Philadelphia: John Benjamins.
Dietz, Klaus. in press. 108. Historical word-formation in English. In Peter O. Müller, Ingeborg Ohnheiser, Susan Olsen & Franz Rainer (eds.), *Word-Formation - An International Handbook of the Languages of Europe*. Berlin: de Gruyter Mouton.
Dossena, Marina. 2012. 55 Late Modern English: Semantics and lexicon. In Alexander Bergs & Laurel J. Brinton (eds.), *English Historical Linguistics*, 887–900. Berlin: Mouton de Gruyter.
Eddington, David. 2000. Analogy and the dual-route model of morphology. *Lingua* 110. 281–298.
Eddington, David. 2002. Spanish diminutive formation without rules or constraints. *Linguistics* 40(2). 395–419.
Eddington, David. 2006. Look Ma, No Rules: Applying Skousen's Analogical Approach to Spanish Nominals in -ión. In Grace Wiebe, Gary Libben, Tom Priestly, Ron Smyth & H. S. Wang (eds.), *Phonology, Morphology, and the Empirical Imperative: Papers in Honour of Bruce L. Derwing*, 371–407. Taipei: Crane.
Embick, David & Alec Marantz. 2008. Architecture and Blocking. *Linguistic Inquiry* 39(1). 1–53.
Fabb, Nigel. 1988. English suffixation is constrained only by selectional restrictions. *Natural Language and Linguistic Theory* 6(4). 527–539.
Fruchterman, Thomas M. J. & Edward M. Reingold. 1991. Graph drawing by force-directed placement. *Journal of Software: Practice and Experience*(21). 1129–1164.
Gahl, Susanne & Alan C. L. Yu (eds.) (2006). *Special Issue on Exemplar-based Models in Linguistics* (23).
Giegerich, Heinz. 1999. *Lexical Strata in English: Morphological Causes, Phonological Effects*. Cambridge: CUP.
Guz, Wojciech. 2009. English affixal nominalizations across language registers. *Poznań Studies in Contemporary Linguistics* 45(4). 447–471.

Hayes, Bruce. 1982. Extrametricality and English stress. *Linguistic Inquiry* 13(2). 227–276.
Hay, Jennifer. 2001. Lexical Frequency in Morphology: Is Everything Relative? *Linguistics: An Interdisciplinary Journal of the Language Sciences* 39(6). 1041–1070.
Hay, Jennifer. 2003. *Causes and Consequences of Word Structure*. London: Routledge.
Kastovsky, Dieter (1986). The problem of productivity in word formation. *Linguistics* 24 (3): 585–600.
Keuleers, Emmanuel. 2008. *Memory-Based Learning of Inflectional Morphology*. PhD. dissertation, Antwerpen: University of Antwerp.
Keuleers, Emmanuel, Dominiek Sandra, Walter Daelemans, Steven Gillis, Gert Durieux & Evelyn Marten. 2007. Dutch plural inflection: The exception that proves the analogy. *Cognitive Psychology* 54(4). 283–318.
Kiparsky, Paul. 1982a. From cyclic phonology to lexical phonology. In Harry van der Hulst & Norval Smith (eds.), *The Structure of Phonological Representations*, 131–175. Dordrecht: Foris.
Kiparsky, Paul. 1982b. Lexical morphology and phonology. In In-Seok Yang (ed.), *Linguistics in the Morning Calm: Selected Papers from SICOL*, 3–91. Seoul: Hanshin.
Liberman, Mark Y. & Alan Prince. 1977. On stress and linguistic rhythm. *Linguistic Inquiry* 8. 249–336.
Lindsay, Mark. 2012. Rival suffixes: synonymy, competition, and the emergence of productivity. In Angela Ralli, Geert Booij, Sergio Scalise & Athanasios Karasimos (eds.), *Morphology and the Architecture of Grammar. Proceedings of the 8th International Morphology Meeting*, 192–203. Patras: University of Patras; URL: http://morbo.lingue.unibo.it/mmm.
Lindsay, Mark & Mark Aronoff. to appear. Natural selection in self-organizing morphological systems. In Fabio Montermini, Gilles Boyé & Jesse Tseng (eds.), *Selected Proceedings of the 7th Décembrettes*. München: LINCOM Europa.
Marchand, Hans. 1969. *Categories and Types of Present-Day English Word-Formation*. München: C.H. Beck.
Marcus, Gary F. Stephen Pinker, Michael Ullman, Michelle Hollander, T. J. Rosen, Fei Xu & Harald Clahsen. 1992. *Overregularization in Language Acquisition* (Monographs of the Society for Research in Child Development). Hoboken, New Jersey: Wiley.
Matthews, Clive A. 2013. On the analogical modelling of the English past-tense: A critical assessment. *Lingua* 133. 360–373.
McClelland, James L. & David E. Rumelhart. 1985. On learning the past tenses of English verbs. In James L. McClelland & David E. Rumelhart (eds.), *Parallel Distributed Processing - Explorations in the Microstructure of Cognition*. Volume 2: Psychological and Biological Models, 216–271. Cambridge, MA: MIT Press.
Pinker, Stephen. 1991. Rules of language. *Science* 253. 530–535.
Pinker, Steven & Alan Prince. 1994. Regular and irregular morphology and the status of psychological rules in grammar. In Susan D. Lima, Roberta Corrigan & Gregory K. Iverson (eds.), *The Reality of Linguistic Rules*, 321–351. Amsterdam / Philadelphia: John Benjamins.
Plag, Ingo. 1996. Selectional restrictions in English suffixation revisited: a reply to Fabb (1988). *Linguistics* 34(4).

Plag, Ingo. 1999. *Morphological Productivity. Structural Constraints in English Derivation*. Berlin: Mouton de Gruyter.
Plag, Ingo. 2003. *Word-Formation in English*. Cambridge: CUP.
Plag, Ingo. 2006. Productivity. In Bas Aarts & April M. McMahon (eds.), *The Handbook of English Linguistics*, 537–556. Oxford: Blackwell.
Prasada, Sandeep & Stephen Pinker. 1993. Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* 8. 1–56.
Raffelsiefen, Renate. 1999. Phonological constraints on English word formation. In Geert Booij & Jaap van Marle (eds.), *Yearbook of Morphology 1998*, 225–287. Dordrecht: Kluwer.
Riddle, Elizabeth M. 1985. A historical perspective on the productivity of the suffixes -ness and -ity. In Jacek Fisiak (ed.), *Historical Semantics - Historical Word-Formation* (Trends in linguistics. Studies and monographs 29), 435–461. Berlin et al.: Mouton de Gruyter.
Romaine, Suzan. 1983. On the productivity of word formation rules and limits of variability in the lexicon. *Australian Journal of Linguistics* 3(2). 177–200.
Säily, Tanja. 2011. Variation in morphological productivity in the BNC: Sociolinguistic and methodological considerations. *Corpus Linguistics and Linguistic Theory* 7(1). 119–141.
Säily, Tanja & Jukka Suomela. 2009. Comparing type counts: The case of women, men and -ity in early English letters. In Antoinette Renouf & Andrew Kehoe (eds.), *Corpus Linguistics: Refinements and Reassessments*, 87–109. Amsterdam: Rodopi.
Skousen, Royal. 1989. *Analogical Modeling of Language*. Dordrecht: Kluwer.
Skousen, Royal. 1992. *Analogy and Structure*. Dordrecht: Kluwer.
Skousen, Royal. 1995. Analogy: A non-rule alternative to neural networks. *Rivista di Linguistica* 7(2). 213–231.
Skousen, Royal. 2002a. An overview of Analogical Modeling. In Royal Skousen, Deryle Lonsdale & Dilworth B. Parkinson (eds.), *Analogical Modeling*, 11–26. Amsterdam / Philadelphia: John Benjamins.
Skousen, Royal. 2002b. Issues in Analogical Modeling. In Royal Skousen, Deryle Lonsdale & Dilworth B. Parkinson (eds.), *Analogical Modeling*, 27–48. Amsterdam / Philadelphia: John Benjamins.
Skousen, Royal. 2005. Analogical Modeling (49). In Reinhard Köhler, Gabriel Altmann & Rajmund G. Piotrowski (eds.), *Quantitative Linguistik*: *Ein internationales Handbuch = Quantitative linguistics* (Handbücher zur Sprach- und Kommunikationswissenschaft = Handbooks of linguistics and communication science = Manuels de linguistique et des sciences de communication / mitbegr. von Gerold Ungeheuer. Hrsg. von Armin Burkhardt … ; Bd. 27), 705–716. Berlin: de Gruyter.
Skousen, Royal. 2009. Expanding Analogical Modeling into a general theory of language prediction. In James P. Blevins & Juliette Blevins (eds.), *Analogy in Grammar*, 164–184. Oxford: OUP.
Skousen, Royal, Deryle Lonsdale & Dilworth B. Parkinson (eds.) (2002). *Analogical Modeling*. Amsterdam / Philadelphia: John Benjamins.
Skousen, Royal & Thereon Stanford. 2007. *AM: Parallel* : available from http://humanities.byu.edu/am/.
van den Bosch, Antal & Walter Daelemans. 2013. Implicit schemata and categories in Memory-based Language Processing. *Language and Speech*. DOI: 10.1177/0023830913484902.

van Tieken-Boon Ostade, Ingrid. 2009. *An Introduction to Late Modern English*. Edinburgh: Edinburgh University Press.
Wedel, Andrew B. 2006. Exemplar models, evolution and language change. *The Linguistic Review* 23. 247–274.
Zamma, Hideki. 2012. *Patterns and Categories in English Suffixation and Stress Placement: A Theoretical and Quantitative Study*: *PhD Dissertation* : University of Tsukuba.