Language Variation and Change, **28** (2016), 1–29.

doi:10.1017/S0954394515000198

Toward more accountability: Modeling ternary genitive variation in Late Modern English

BENEDIKT SZMRECSANYI
KU Leuven

DOUGLAS BIBER
Northern Arizona University

JESSE EGBERT
Brigham Young University

KARLIEN FRANCO
KU Leuven

ABSTRACT

Whereas the alternation between the *s*-genitive (*the New Year's message*) and the *of*-genitive (*the message of the New Year*) is well documented, our study offers a more accountable analysis of genitive variation by including noun-noun (NN)-genitives (*the New Year message*). We consider four different variable contexts (*s* versus *of*, NN versus *of*, NN versus *s*, and NN versus *s* versus *of*), which we analyze using regression analysis. The dataset consists of 10,054 variable genitives drawn from the *Representative Corpus of Historical English Registers*. The material covers the period between 1650 and 2000, thus enabling us to track the evolution of variable genitive grammar in real time. We report that there is an overall drift toward the NN-genitive, which is preferred over other variants when constituent noun phrases are short, possessor constituents are inanimate, and possessum constituents are thematic. In addition to these substantive contributions, we showcase methods of dealing with a complex dataset covering nonbinary grammatical variation.

This paper is concerned with variation and change in the expression of genitive relations in the Late Modern English (henceforth: LateModE) period, which in this study we take to last from the middle of the 17th century to the end of the 20th century. Our primary goal is to challenge the "consensus" (Rosenbach, 2014:222) in most previous research that the English genitive has only two

Thanks go to Anette Rosenbach, Dirk Speelman, Jason Grafmiller, Benedikt Heller, and two anonymous referees for excellent comments and feedback. The first author gratefully acknowledges an Odysseus grant by the Research Foundation Flanders (FWO, grant no. G.0C59.13N). The usual disclaimers apply.

variants: the *s*-genitive and the *of*-genitive. What we hope to demonstrate instead is that there are no less than three syntactic variants. Consider examples (1) to (3):

(1) **the *s*-genitive**:
Mr Floyd Clark, the [***FBI***]$_{\text{possessor}}$***'s*** [***director of criminal investigations***]$_{\text{possessum}}$, *said he had leads in the investigation* . . . (1989tim2.n8b)[1]

(2) **the *of*-genitive**:
. . . the problems raised by the Pope apply equally to other groups with a grievance such as the Armenians, the [***Catholics***]$_{\text{possessum}}$ ***of*** [***Northern Ireland***]$_{\text{possessor}}$ *and the Palestinians*. (1989tim2.n8b)

(3) **the NN-genitive**:
The Pope, in his [***New Year***]$_{\text{possessor}}$ [***message***]$_{\text{possessum}}$, *denounced terrorism as a way of supporting the causes of minorities.* (1989tim2.n8b)

The *s*-genitive in (1) can be broadly paraphrased as *the director of criminal investigations of the FBI* or *the FBI director of criminal investigations*. The *of*-genitive in (2) translates into *Northern Ireland's Catholics* or *Northern Ireland Catholics*; and the noun-noun (henceforth: NN) genitive[2] in (3) can be alternatively expressed as either *his New Year's message* or *his message of the New Year*.

Examples (1) to (3) are two-way interchangeable (i.e., all three variants are possible), but there are many *s-, of-*, and NN-genitives that are interchangeable with only one of the other genitive variants, but not both. Consider, for example, the NN-genitive in *its trace can be followed as a darker zone through the wheat fields* (1975tcha.s8b): *Wheat fields* can be paraphrased as *fields of wheat*, but not as ?*wheat's fields*. The fact that not all genitive variants in our dataset can be paraphrased by all other genitive variants poses analytical challenges, which we address as follows.

Using variationist methodology (e.g., Jankowski & Tagliamonte, 2014; Labov, 1982), we analyze the news, science, and letters section of the ARCHER corpus, from which we extract 10,054 *s*-genitives, *of*-genitives, and NN-genitives that are interchangeable with at least one of the other genitive options. We annotate the genitives in this dataset for the extent to which they are interchangeable with one or both of the other variants, as well as for a range of predictors[3] known to constrain genitive variation, such as possessor animacy and constituent length. Subsequently, we slice up the dataset into four different variable contexts: three binary alternation contexts (interchangeable *s*-genitives versus *of*-genitives, NN-genitives versus *of*-genitives, and NN-genitives versus *s*-genitives), plus one ternary alternation context where each variant is paraphrasable by *both* of the other variants. We then analyze these four alternation contexts one-by-one using regression analysis.

Why do we need this study? Variation between the *s*-genitive and the *of*-genitive is extremely well researched (the literature is too voluminous to be reviewed here in any detail; see Rosenbach [2014] for a comprehensive literature review). Suffice it to say that this body of research generally concludes that the *s*-genitive has been on the rise, vis-à-vis the *of*-genitive, since the Early Modern English period (and is spreading right now, particularly in informational registers such as newspaper prose). Several of these studies have also found that animate and short possessors in particular favor the *s*-genitive, whereas short possessums and possessors ending in a final sibilant favor the *of*-genitive.

But much less is known about the patterns of variation between the *s*-genitive and the *of*-genitive when the NN-genitive is included as a possible variant. Rather, most researchers have tended to ignore the NN-genitive, simplistically restricting attention to the binary variation between the *s*-genitive and the *of*-genitive. Studies that actually acknowledge the existence of the NN-genitive (Feist, 2012; Rosenbach, 2006, 2007a, 2007b) or discuss NN constructions as alternatives to *of*-genitives (Biber & Gray 2011, 2013, 2016) tended to eschew a quantitative-variationist perspective. Our theoretical/descriptive objective here is to redress this negligence: how does the choice between the NN-genitive and either the *s*-genitive or the *of*-genitive differ from the well-known alternation between the *s*-genitive and the *of*-genitive? What are the relevant predictors, and how did they evolve during the LateModE period? Considering all three genitive variants, how does the system work?

On the methodological plane, the present study explores the techniques required for research that attempts to be compliant with the Principle of Accountability (Labov, 1969:fn 20) in cases involving nonbinary grammatical variation. There is surely enough previous work on linguistic variation encompassing $n > 2$ variants: consider, for example, Sankoff and Rousseau (1989), who extended the variable rule methodology to handle phonological variables with more than two variants. Another approach to multiple variants can be found in research on copula deletion (see, e.g., Baugh, 1980; Labov, 1969; Meyerhoff & Walker, 2007; Poplack & Sankoff, 1987; Rickford, Ball, Blake, Jackson, & Martin, 1991; Wolfram, 1974). In these studies, ratios of one or two of the three variants (full form, contraction, and deletion) to the other variants are calculated (a comparison of the different calculations of contraction and deletion can be found in Rickford et al. [1991]). Note, however, that the research we report in this paper faces different challenges: unlike copula variation phenomena, the genitive variants we study here cannot be arranged on a scale of reduction (or some other process); in addition, our dataset is unusually complex because we are faced with a number of partially overlapping though not entirely congruent variable contexts (recall that not every genitive variant can be paraphrased by every other genitive variant in all contexts).

Thus, a different methodological approach is needed. In the next section, we introduce the corpus used for our analyses, and following that we discuss the statistical methods used to analyze ternary patterns of variation with the complications noted. The ensuing sections then present the main findings regarding the patterns of genitive variation: we survey the development of variant proportions in the course of the LateModE period, present the regression models, and offer a discussion of the findings. Finally, we add some concluding remarks and sketch directions for future research.

DATA

We analyze ARCHER, *A Representative Corpus of Historical English Registers* (Biber, Finegan, & Atkinson, 1994; Yáñez-Bouza, 2011). ARCHER covers the period between 1650 and 1999, with the precise year of composition for each text being typically available. We restrict attention to British texts and three particular registers: science (151,786 words of running text, 70 texts), letters (89,888 words of running text, 187 texts), and news (159,054 words of running text, 70 texts). The rationale is as follows: both science and news are good representatives of the written mode, with science writing being especially innovative in the use of phrasal complexity features (Biber & Gray, 2016). By contrast, letters—in particular, private letters—are considered fairly oral in nature (Raumolin-Brunberg, 2005:57).

METHODS

Defining the variable contexts and variant extraction

We consider three genitive variants: the *s*-genitive, as in (1); the *of*-genitive, as in (2); and the NN-genitive, as in (3) (as we will mention in the Conclusion, there may be additional variants, which we do not, however, include in the present analysis). We inspected the corpus material to (i) identify occurrences of *s*-genitives, *of*-genitives, and premodifying nouns, and (ii) determine their interchangeability with the other two variants. For *s*-genitives and *of*-genitives, we followed the guidelines in Hinrichs and Szmrecsanyi (2007) and Wolk, Bresnan, Rosenbach, and Szmrecsanyi (2013), and we then developed a similar set of methods for coding NN-genitives.

The first step was to identify potentially relevant occurrences of each of the three variants. We began by automatically tagging all occurrences of *of* and final **'s/*s'* (as well as final **s* in 17th- to 18th-century texts, because apostrophes were often omitted from *s*-genitives during that period). For the NN-genitives, we used the Biber Tagger to automatically identify all nouns, and then searched for instances of two adjacent nouns.

We then manually coded each occurrence, to first eliminate cases that were not actually *s*-genitives, *of*-genitives, or NN-genitives, and to then mark, in the remaining cases, the boundaries of the two noun phrases (NPs) and their syntactic heads. Once again, we refer to genitive constituents as the "possessor" and the "possessum," even though many occurrences do not actually express the meaning of possession.

The second step was to analyze each genitive construction by hand to determine whether it was "interchangeable" with one or both of the other two variants; we were thus interested in the extent to which we are dealing with "alternate ways of saying 'the same' thing" (Labov 1972:188). This step was based on our intuitions, deciding whether the structure could be rephrased with another variant to express essentially the same meaning.[4] We required that the rephrasing use the same lexemes, with two exceptions: the optional addition or deletion of a determiner to the possessum for *of*-genitives (e.g., *the government's policy* versus *the policy of the government*), and the optional pluralization or singularization of the possessor for premodifying nouns (e.g., *home prices* versus *prices of homes*). The following special cases were coded as not interchangeable: (i) phrases that have been conventionalized (e.g., *Murphy's law; post office*); (ii) constructions in which an *s*-genitive is not followed by an explicit possessum phrase (e.g., *Her memory is like an elephant's*)—paraphrasing these constructions would require the addition of lexical items or would result in a phrase introduced by a different preposition than *of*, which is why such genitives are customarily excluded in the literature on genitive variation (for discussion, see Hinrichs and Szmrecsanyi [2007:446]); (iii) titles of books, films, or other works that are premodified with an *s*-genitive (e.g., *Van Gogh's Starry Night*); (iv) measures expressed as *of*-genitives (e.g., *three gallons of milk*); (v) *of*-genitives where the possessor NP has a postmodifier (e.g., *the girlfriend of the man that I met*); and (vi) noun premodifiers that are not definite (because the possessum in *s*-genitives is always definite—e.g., *a London college*).

We note explicitly that the dataset includes the following genitive types: (i) determiner *s*-genitives (*Einstein's paper*), determiner nominal premodifiers (*the Einstein paper*), and (complement) *of*-genitives (*the head of Einstein*); (ii) classifying *s*-genitives (*driver's licence*) as well as parallel nominal premodifiers (*driver licence*), unless the phrases are conventionalized; and (iii) objective genitives of the type *the release of the prisoner/the prisoner's release/the prisoner release.*

Prepositional phrases not headed by *of* (e.g., *license for drivers*) are categorically excluded. We should also emphasize that we sometimes annotated NN-genitives as interchangeable with *of* even though they may seem much more natural with another preposition. For example, *union negotiators* may be paraphrased as *negotiators of unions*, which is why we annotated this NN-genitive as being interchangeable with the *of*-genitive—but that said, the more natural option is certainly *negotiators from unions*. It moreover bears mentioning that we included a slew of different semantic relations—in essence, anything that was interchangeable with either an *s*-genitive or an *of*-genitive. But we do explicitly acknowledge that certain sematic relations are probably more likely to be interchangeable than others, resulting in potential interaction between interchangeability (and so inclusion as a "genitive"), semantic relation, determiner versus classifier status, and register.

By way of illustration, Table 1 contains examples of interchangeable genitives for each of the possible variants. Beyond the guidelines enumerated here, coders were instructed to rely on their best judgment to determine interchangeability based on whether it could equally (or plausibly) be expressed as the form in question by applying a simple conversion rule. After several rounds of trial coding and subsequent revisions to the coding scheme, two coders rated a sample in order to measure intercoder reliability. Reliability was calculated for each of the three genitive variants using simple percentage agreement and

TABLE 1. *A typology of interchangeable genitives*

Variant	Interchangeable With	Example
NN-genitive	*s*-genitive not *of*-genitive	*satellite photographs* (1967stm2.n8b)
	of-genitive not *s*-genitive	*the insertion point* (1975macm.s8b)
	s-genitive and *of*-genitive	*the Castro forces* (1959man1.n8b)
of-genitive	*s*-genitive not NN-genitive	*the information of the people* (1665int2.n2b)
	NN-genitive not *s*-genitive	*samples of data* (1975macm.s8b)
	s-genitive and NN-genitive	*the lords of the council* (1654mer2.n2b)
s-genitive	*of*-genitive not NN-genitive	*his mother's death* (1701trot.d3b)
	NN-genitive not *of*-genitive	*Sunday's post* (1815aust.x5b)
	of-genitive and NN-genitive	*St. James's coffee-house* (1710adds.x3b)

Cohen's *k*. For these calculations, in order to achieve agreement for a given genitive, the two coders needed to agree on a single classification for the genitive from the following four possible categories: (i) noninterchangeable, (ii) interchangeable with variant 1 only, (iii) interchangeable with variant 2 only, (iv) interchangeable with both variant 1 and variant 2. The *s*-genitive ($n = 84$) coding achieved a simple agreement rate of 95% and a "very good" Cohen's *k* of .91. Reliability for *of*-genitives ($n = 112$) achieved a simple agreement of 90% and a "very good" Cohen's *k* of .80. Finally, the reliability analysis for NN-genitives ($n = 91$) yielded a lower, yet still acceptable simple percentage agreement of 85%, with a "good" Cohen's *k* of .69.

In total, the dataset subject to analysis in this paper covers 10,054 interchangeable genitives.

ANNOTATION

Relevant genitive occurrences were annotated for a range of predictors, mostly the usual suspects from the literature on variation between the *s*-genitive and the *of*-genitive. With the exception of possessor animacy, all predictors were coded automatically using computer programs.

Constituent length of the possessor (POR_LENGTH_WORDS) and the possessum (PUM_LENGTH_WORDS)

In languages such as English, longer, heavier elements tend to follow shorter, lighter ones. This is known as the principle of end weight (Behaghel, 1909; Biber, Johansson, Leech, Conrad, & Finegan, 1999:898; Wasow, 2002), which can in turn be subsumed under a more general principle. MacDonald (2013) argued that language users' planning and production of utterances is guided by a number of biases, one of which is the "Easy First" principle: when faced with a word order choice, language users will put "easier" elements first, to maximize planning time for more difficult elements; "easy" MacDonald (2013:3) defined as "more frequent, shorter . . . , less syntactically complex, more important or conceptually salient to the speaker, and previously mentioned ('given') in the discourse." Hence, for example, if the possessum is short, the *of*-genitive should be preferred over the *s*-genitive and the NN-genitive because the *of*-genitive places the possessum before the possessor. Analyzing a dataset similar to ours, Ehret, Wolk, and Szmrecsanyi (2014:297) showed that using separate measures for possessor and possessum length is superior to other ways to model weight effects, and so we likewise use separate measures. For the sake of simplicity, we moreover abstain from combining the two length predictors with ratio measures, or length difference measures. As for the actual measurements, we approximate heaviness by counting the number of orthographic words, again relying on the Ehret et al. (2014) study for justification. Consider the genitive phrase in (1) (*FBI's director of criminal investigations*): the possessor phrase, *FBI*, has a length (weight) of one word, and the possessum phrase, *director of criminal investigations*, has a length of four orthographic words. Note that the counts do not include determiners, as including these results in various complications (for example, *s*-genitive possessors never have a definite determiner).

Possessor animacy (ANIMACY)

Possessor animacy is considered one of the most crucial factors constraining variation between the *s*-genitive and the *of*-genitive (Jankowski & Tagliamonte, 2014; Rosenbach, 2005): animate possessors attract the *s*-genitive. This preference is compatible with the Easy First principle, as animate nouns are known to be more easily retrievable from memory (MacDonald, 2013:5). Crucially, there is reason to believe that animacy also plays a role in the choice between the NN-genitive and other options (see Rosenbach, 2006:90, 2007a:177, 2007b:166). We thus hand-coded possessor heads for animacy, assigning the value "animate" if the possessor head noun belonged to the categories HUMAN or ANIMAL in the coding scheme of Zaenen, Carlette, Garretson, Bresnan, Koontz-Garboden, Nikitina, O'Connor, and Wasow (2004), and "nonanimate" otherwise. For example, the possessor in (1), *FBI*, is classed as nonanimate.

Alpha-persistence (ALPHA_PERSISTENCE_OF, ALPHA_PERSISTENCE_S, ALPHA_PERSISTENCE_NN)

This predictor set is about persistence, also known as priming (Gries, 2005; Travis, 2007) or the "plan reuse" principle (MacDonald, 2013), another bias in planning and production. More specifically, we are interested here in what Szmrecsanyi (2006:3–5) called "alpha persistence": the tendency of language users to reuse syntactic structures in choice contexts that they had used in previous choice contexts. The predictors thus indicate whether the genitive variant in question had been used at least once (but possibly more than once) in a stretch of 50 words prior to the genitive slot under analysis (note that the prime had to be interchangeable as well). The prediction is that recent usage of, say, a variable *s*-genitive makes reusage of the *s*-genitive more likely.

Beta-persistence (BETA_PERSISTENCE_OF, BETA_PERSISTENCE_S, BETA_PERSISTENCE_NN)

Here we are concerned with what Szmrecsanyi (2006:3–5) called “beta persistence”: the tendency of language users to select variants in choice contexts that resemble in some way recently used linguistic material that is not necessarily itself variable. So, BETA_PERSISTENCE_OF checks whether the preposition *of* (as in *of course*) had been used at least once (but possibly more than once) in a stretch of 50 words prior to the genitive slot under analysis; BETA_PERSISTENCE_S checks for previous usage of clitic *s*; and BETA_PERSISTENCE_NN checks for previous usage of a generic NN string. Again, we hypothesize that recent usage of similar material makes usage of a particular genitive variant more likely.

Type-token ratio (TTR)

Hinrichs and Szmrecsanyi (2007) calculated type-token ratios as a proxy for lexical density, finding that lexically dense contexts favor the *s*-genitive. In like manner, we calculate type-token ratios of the textual passages where the genitive observations are embedded (i.e., starting 50 words before and ending 50 words after a given genitive construction).

Thematicity of the possessor (POR_THEMATICITY) and the possessum (PUM_THEMATICITY)

Osselton (1988) claimed that the general topic of a text partly determines which nouns in that text can take the *s*-genitive. Hinrichs and Szmrecsanyi (2007) confirmed this claim, operationalizing thematicity by establishing the normalized text frequency of the possessor NP’s head noun in the corpus text in which the genitive construction occurs. We proceed analogously and additionally establish possessum thematicity by determining the normalized text frequency of the possessum NP’s head noun. For example, the possessor head noun of the genitive construction in (1) (*FBI*) has a local text frequency of 1.91 occurrences per 1000 words (ptw) of running text; the corresponding value for the possessum head noun (*director*) is 1.72. We are, to our knowledge, the first to explore possessum thematicity effects. The general prediction is that more thematic constituents will be placed in first position. This expectation once again follows from the Easy First principle, in that increased thematicity translates into increased frequency, salience, and, potentially, previous mention (MacDonald, 2013:3), factors which all should increase ease of retrieval.

Final sibilancy (FINAL_SIBILANCY)

A final sibilant in the possessor discourages usage of the *s*-genitive (e.g., Zwicky, 1987). Thus we used a string matching algorithm to identify all possessors ending in <s> (as in *Congress*), <z> (as in *jazz*), <ce> (as in *resistance*), <sh> (as in *Bush*), <tch> (as in *match*), or <dge> (as in *judge*).

Lexical effects (POR_HEAD_NOUN AND PUM_HEAD_NOUN)

To model by-item effects, we broadly follow the method in Wolk et al. (2013) and consider both the lemma of the possessor head noun and of the possessum head noun as random effects in regression analysis. Also as in Wolk et al. (2013), we collapsed all head nouns that did not reach a threshold of at least four observations into an "other" category. The rationale was (i) to facilitate computation of the models by reducing the number of levels to be considered and (ii) to avoid making the random effect structure too sensitive to particularities.

Metadata and language-external variables

FILENAME approximates author idiosyncrasies via a random effect. As modality and register are known to play a role in genitive choice (Grafmiller, 2014), REGISTER models differences among the news, science, and letters registers considered in the present study, also via a random effect.[5] We use the fixed-effect predictor TIME to investigate probabilistic grammar change by exploring two-way interactions between TIME and the language-internal predictors. Note that in the regression models, we model real time by dividing the year of creation of a text by 100, which yields centuries, and subsequently centering the variable around the year 1800 (see Wolk et al., 2013).

SLICING UP THE DATASET AND STATISTICAL ANALYSIS

We rely on binary logistic regression analysis to model the variation at hand. Regression analysis has been customary in variation studies since the 1970s (Sankoff & Labov, 1979); we will use a modern refinement of the technique, mixed-effects models (Pinheiro & Bates, 2000), as implemented in the R library lme4 (Bates, Maechler, & Bolker, 2012).[6] Ideally, the analyst will want to fit one unitary model to the dataset. This is not possible here, as not each genitive in our dataset is paraphrasable by *both* other genitive variants. Consider once again the NN-genitive in *its trace can be followed as a darker zone through the **wheat fields*** (1975tcha.s8b): *Wheat fields* can be paraphrased as *fields of wheat*, but not as ?*wheat's fields*. It is because of this complication that fitting a unitary model to the dataset would violate the Principle of Accountability (Labov, 1969: fn 20), which is why we slice up the dataset into four different variable contexts:

1. ***s*-genitives and *of*-genitives** that are interchangeable with each other, but not with NN-genitives (model 1): $n = 4195$ ($n_{\text{s-genitive}} = 831$, $n_{\text{of-genitive}} = 3364$)
2. **NN-genitives and *of*-genitives** that are interchangeable with each other, but not with *s*-genitives (model 2): $n = 2832$ ($n_{\text{NN-genitive}} = 905$, $n_{\text{of-genitive}} = 1927$)
3. **NN-genitives and *s*-genitives** that are interchangeable with each other (model 3): $n = 676$ ($n_{\text{NN-genitive}} = 563$, $n_{\text{s-genitive}} = 113$)
4. **NN-genitives versus *of*-genitives versus *s*-genitives** – all occurrences that are interchangeable with *all* other variants (model 4): $n = 2927$ ($n_{\text{NN-genitive}} = 470$, $n_{\text{of-genitive}} = 2351$, $n_{\text{s-genitive}} = 106$)

We will present one regression model for each of the four contexts. How did we arrive at these models? The issue of model fitting and model selection is a matter of current debate in the literature on regression analysis. We report here so-called *minimal adequate models*, from which predictors that "do not contribute enough to the model's success" (Gries, 2013:260) have been eliminated.[7] To trim down the models, we broadly followed the top-down strategy described in Zuur, Ieno, Walker, Saveliev, and Smith (2009:121–122): we first optimized the random-effects structure in Model 1, finding that the random effect candidates FILENAME, REGISTER, POR_HEAD_NOUN, and PUM_HEAD_NOUN all are, as adjustments to the intercept, warranted statistically according to likelihood ratio tests. For the sake of comparability, we kept this random-effect structure constant throughout the rest of the study, without reoptimizing in Models 2, 3, and 4. Next, for each of the four models we fitted the maximal models including, as main effects, all language-internal predictors and TIME as well as interaction terms between TIME and the language-internal predictors.[8] Subsequently, the models were simplified (Gries, 2013:260, calls this "backward selection") by removing factors and interaction terms lacking significant explanatory power, starting with the least significant higher-order interaction, and moving toward the main effects. Note here that our approach to model selection is largely compatible with Gelman and Hill's (2007:69) guidelines, except that we consistently excluded predictors that are not statistically significant, even if the effect direction was the expected one.

Note that even with guidelines as sketched herein, model fitting to some extent remains an inevitably subjective endeavor. To address this subjectivity, two fitters—the first-named and the last-named author—independently pruned the maximal models, which enables us to report qualitative measures of "interfitter reliability."

The dataset and a commented version of the R code necessary to replicate our analysis are provided in the online appendix (the online appendix can be viewed at http://dx.doi.org/10.1017/S0954394515000198).

VARIANT FREQUENCIES AND RATES IN REAL TIME

We begin by discussing the development of variant frequencies and rates in real time. Figure 1 presents four graphs—one for each variable context—with each graph plotting variant rates against real time; the data labels in the diagrams indicate raw frequencies, for the sake of giving some idea of the balance of observations over the different time periods. As for the *s*-genitive versus *of*-genitive alternation (top-left plot), there are few surprises given previous research based on a similar dataset (e.g., Wolk et al., 2013). The NN-genitive versus *of*-genitive (top-right plot) and NN-genitive versus *s*-genitive (bottom-left plot) alternations are characterized by fairly drastic frequency shifts in favor of the NN-genitive. By contrast, the ternary alternation (bottom-right plot) appears to be more stable, although we diagnose quite some movement in the 20th century—again, in favor of the NN-genitive. We conclude that in all those

Fig. 1 - B/W online, B/W in print

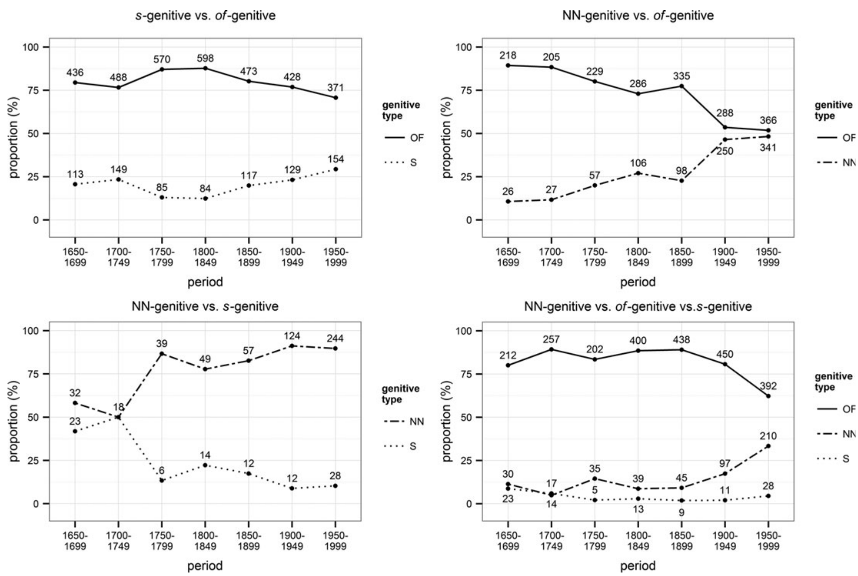

FIGURE 1. Variant proportions (%, *y*-axis) in real time (*x*-axis); data labels indicate raw *n*s (not normalized). Top left plot: *s*-genitive versus *of*-genitive (Model 1). Top right plot: NN-genitive versus *of*-genitive (Model 2). Bottom-left plot: NN-genitive versus *s*-genitive (Model 3). Bottom right plot: ternary variation (Model 4).

datasets where the NN-genitive is involved, it is on the rise. The *of*-genitive, by contrast, is overall on the decline.

MULTIVARIATE ANALYSIS

We now present separate analyses for each of the four variable contexts subject to analysis. Table 2 summarizes the fixed effects in the models, and Table 3 the random effects.

Modeling the choice between the s-genitive and the of-genitive (Model 1)

We begin by exploring “classical” variation between the *s*-genitive, as in (4), and the *of*-genitive, as in (5). Unlike previous studies, we restrict attention to *s*-genitive and *of*-genitive occurrences that are interchangeable with each other but not with the NN-genitive.[9] The dataset on which Model 1 is based covers $n = 4195$ observations ($n_{\text{s-genitive}} = 831$, $n_{\text{of-genitive}} = 3364$).

(4) **the *s*-genitive**:
Next [***year***]$_{\text{possessor}}$ ***’s*** [***reshuffle***]$_{\text{possessum}}$ *may be accompanied by a reorganization of Whitehall, according to senior government sources.* (1989tim2.n8b)

TABLE 2. *Regression coefficients (b, log odds) in mixed-effects binary logistic regression; fixed effect estimates only*

	s-genitive vs. *of*-genitive (Model 1)				NN-genitive vs. *of*-genitive (Model 2)				NN-genitive vs. *s*-genitive (Model 3)				NN-genitive vs non-NN-genitive (Model 4)			
Predicted outcome	*s*-genitive				NN-genitive				NN-genitive				NN-genitive			
n	831/4195				905/2832				563/676				470/2927			
Parameters	*b*	sig.	*n*	% *s*-genitive	*b*	sig.	*n*	% NN-genitive	*b*	sig.	*n*	% NN-genitive	*b*	sig	*n*	% NN-genitive
Intercept	−3.77	***			−2.23	***			4.32	***			−4.17	***		
Constituent length (1 unit corresponds to 1 orthographically transcribed word)																
POR_LENGTH_WORDS	−.77	***	continuous variable		−.58	***	continuous variable				continuous variable		−.81	***	continuous variable	
PUM_LENGTH_WORDS	.16	*	continuous variable		−3.16	***	continuous variable		−.70	**	continuous variable		−1.86	***	continuous variable	
Possessor animacy																
ANIMACY = "inanimate"	default level		2799	7	default level		2711	33	default level		585	88	default level		2727	16
ANIMACY = "animate"	3.31	***	1396	45			121	18	−2.48	***	91	52			200	18
Alpha-persistence																
ALPHA_PERSISTENCE_OF = "no"	default level		1595	23	default level		945	28	default level		168	82	default level		768	14
ALPHA_PERSISTENCE_OF = "yes"			2600	18			1887	34			508	84	0.41	*	2159	17
ALPHA_PERSISTENCE_S = "no"	default level		3541	19	default level		2427	32	default level		576	84	default level		2573	16
ALPHA_PERSISTENCE_S = "yes"	−1.06	***	654	26			405	34			100	79	(−.28)		354	19
ALPHA_PERSISTENCE_NN = "no"	default level		3418	19	default level		2032	27	default level		448	80	default level		2188	14
ALPHA_PERSISTENCE_NN = "yes"			777	22			800	45			228	89			739	23
Beta-persistence																
BETA_PERSISTENCE_OF = "no"	default level		641	28	default level		414	31	default level		75	85	default level		305	18
BETA_PERSISTENCE_OF = "yes"			3554	18			2418	32			601	83			2622	16
BETA_PERSISTENCE_S = "no"	default level		3593	18	default level		2489	32	default level		571	85	default level		2631	15
BETA_PERSISTENCE_S = "yes"	1.15	***	602	32			343	31			105	75			296	24
BETA_PERSISTENCE_NN = "no"	default level		1719	18	default level		995	21	default level		202	78	default level		1192	11
BETA_PERSISTENCE_NN = "yes"			2476	21	.60	***	1837	38			474	86	.45	*	1735	20

Thematicity (1 unit corresponds to 1 ptw point)																
POR_THEMATICITY			continuous variable				continuous variable				continuous variable		−.16	**	continuous variable	
PUM_THEMATICITY	.11	**	continuous variable		.23	***	continuous variable		(.09)		continuous variable		.20	***	continuous variable	
Final sibilancy																
FINAL_SIBILANCY = "no"	default level		3055	24	default level		2118	37	default level		597	84	default level		2205	19
FINAL_SIBILANCY = "yes"	−1.69	***	1140	8	−2.27	***	714	18	(−.58)		79	77	−1.61	***	722	8
Real TIME (1 TIME unit corresponds to 1 century since 1800)																
TIME	.80	***			.87	***			.61	***			.78	***		
Interactions involving TIME																
PUM_LENGTH_WORDS * TIME			interaction				interaction				interaction		−.75	**	interaction	
ANIMACY = "animate" * TIME	−.52	***	interaction				interaction				interaction				interaction	
ALPHA_PERSISTENCE_S = "yes" * TIME			interaction				interaction				interaction		.88	**	interaction	
POR_THEMATICITY * TIME			interaction				interaction				interaction		.16	***	interaction	
PUM_THEMATICITY * TIME	−.07	*	interaction				interaction		.13	*	interaction		.06	*	interaction	
FINAL_SIBILANCY = "yes" * TIME	−.52	**	interaction		.63	*	interaction		.89	*	interaction				interaction	
Model quality measures																
C	.96				.96				.94				.96			
% correctly predicted (% baseline)	91.8% (80.2%)				89.2% (68.0%)				90.7% (83.3%)				93.0% (84.0%)			

Note: Negative coefficients disfavor the predicted outcome (i.e., the first-listed variant in each column), positive coefficients favor the predicted outcome. Empty cells in the *b* column indicate that the predictor is not part of the minimal adequate model. Bracketed coefficients indicate nonsignificant main effects that are included in the minimal adequate model because of significant interactions. Significance codes: *** $p < .001$; ** $p < .01$; * $p < .05$. POR = possessor, ptw = per thousand words, PUM = possessum.

TABLE 3. *Variances associated with random effects (intercept adjustments) in mixed-effects binary logistic regression; bigger values indicate greater importance of the random effect*

	s-genitive vs. *of*-genitive (Model 1)	NN-genitive vs. *of*-genitive (Model 2)	NN-genitive vs. *s*-genitive (Model 3)	NN-genitive vs. non-NN-genitive (Model 4)
POR_HEAD_NOUN	2.39	7.92	13.04	5.76
PUM_HEAD_NOUN	1.86	3.64	0.0	4.90
FILENAME	0.79	0.35	0.85	1.02
REGISTER	0.09	0.21	0.00	0.34

(5) **the *of*-genitive**:
The sharp [***intensification***]$_{\text{possessum}}$ **of** [***the conflict***]$_{\text{possessor}}$ *contrasts with a claim by Bishop Abel Muzorewa . . .* (1979stm1.n8b)

Finding the minimal adequate model was unproblematic: the two fitters independently arrived at the exact same minimal adequate model. The index of concordance *C*, which measures how well Model 1 discriminates between genitive outcomes, is an excellent .96; note here that values > .8 are interpreted as indicating a good fit (Tagliamonte & Baayen, 2012:156). As for predictive accuracy, the model correctly predicts 92% of all outcomes vis-à-vis the baseline prediction of 80% (80% of all genitives in this dataset are *of*-genitives). Multicollinearity is not a problem, as Model 1's condition number ($\kappa = 3.5$) is below the customary threshold of 15.[10] The fixed-effect estimates discussed here are stable under bootstrap validation (Baayen, 2008:283).

We move on to a brief discussion of the four random effects—which are all adjustments to the intercept—in Model 1. Table 3 reports variances associated with the random effects in the models; their magnitude is proportional to relative importance. So in Model 1, the two lexical by-item effects, PUM_HEAD_NOUN and, in particular, POR_HEAD_NOUN, do most of the heavy lifting. As for lexical effects in the possessum, inspection of the intercept adjustments reveals that, for example, possessum *myrrh* (as in *the myrrh of the ancients*, 1775bruc.s4b) disfavors the *s*-genitive while possessum *news* (as in *this evening's news*, 1866pal1.n6b) strongly favors the *s*-genitive, all other things being equal. Turning to the possessor phrase, we note that, for example, possessor *men* (as in *the rescue of the men*, 1919dai2.n7b) disfavors the *s*-genitive, whereas possessor *Friday* (as in *Friday's rush*, 1928tim1.n7b) strongly favors the *s*-genitive. There is also quite a bit of by-subject variation (FILENAME)—for example, the author of text 1819mor2.n5b disfavors the *s*-genitive, all other things being equal, while the author of text 1751rich.x4b seems to have a soft spot for the *s*-genitive. Finally, we also see REGISTER variation: science texts receive negative intercept adjustments, whereas news texts and especially letters receive positive adjustments.

Let us now have a look at the fixed-effects structure of Model 1. Table 2 reports the regression coefficients. Positive coefficients indicate that a particular level of a categorical predictor variable, vis-à-vis the default level, or a one-unit increase in a continuous predictor variable, make a particular linguistic outcome—in this case: choice of an *s*-genitive rather than an *of*-genitive—more likely. Negative coefficients indicate that the predicted outcome is less likely under the conditions specified. The absolute value of the coefficient is proportional to the magnitude of the effect. Thus, for example, the predictor variable POR_LENGTH_WORDS has a coefficient of −.77 in Model 1. Thus longer possessor phrases disfavor the *s*-genitive, exactly as they should as per the principle of end weight. To make regression coefficients more interpretable, we may convert them into so-called odds ratios by raising *e* to the power of the coefficient: $e^{-.77} = .46$. Hence, for every one-word increase in the length of the possessor phrase, the odds for an s-genitive are reduced by a factor of .46.

By the same logic, longer possessums (PUM_LENGTH_WORDS) favor the *s*-genitive. According to the coefficient associated with ANIMACY = "animate," animate possessors hugely favor the s-genitive: if the possessor head noun is animate instead of inanimate, the odds for the s-genitive increase by a factor of $e^{3.31} = 27.4$. Strong animacy effects were to be expected (Rosenbach, 2005). By contrast, the negative coefficient associated with ALPHA_PERSISTENCE_S = "yes" is a bit surprising, as one would have expected that recent usage of an interchangeable *s*-genitive favors reuse of the *s*-genitive, which it does not. BETA_PERSISTENCE_S = "yes" does have the theoretically expected effect: recent usage of an *s*-clitic favors usage of the *s*-genitive. As for the thematicity predictors, previous research (Hinrichs & Szmrecsanyi, 2007) found significant possessor thematicity effects, but in our dataset the predictor is not in the minimal adequate model due to lacking significance. Instead, possess*um* thematicity—recall that our study is the first to explore thematicity of the possessum—has a significant main effect: the positive coefficient of PUM_THEMATICITY shows that more frequent possessum heads tend to attract the *s*-genitive. This effect direction is interesting: as per the Easy First principle (MacDonald, 2013), one would have predicted that textually frequent (i.e., accessible and hence "easy") possessum heads are preferentially coded with the *of*-genitive (which would place the possessum first), which they are not. FINAL_SIBILANCY = "yes" has the expected sign: a final sibilant in the possessor discourages usage of the *s*-genitive, as reported in the literature. The positive coefficient associated with the predictor TIME is no surprise either, as we know that the s-genitive is diachronically on the rise (see, e.g., Wolk et al., 2013). Model 1 suggests that for every passing century in the LateModE period, the odds for the s-genitive increase by a factor of $e^{.80} = 2.3$.

We also tested for interaction effects between the language-internal predictors and TIME, to determine whether the probabilistic grammar of genitive choice has changed during the LateModE period. In the dataset underlying Model 1, we found three significant interaction effects. First, the interaction term ANIMACY = "animate" * TIME has a negative coefficient (−.52). Because TIME is centered around AD 1800, the main effect coefficient of animacy in the first half of the table describes the effect that ANIMACY = "animate" has in the year AD 1800. To calculate the effect of ANIMACY = "animate" in later centuries, we subtract .52 from the main effect coefficient for each passing century; to calculate the effect of ANIMACY = "animate" in earlier centuries, we add .52 for each passing century. This means that the animacy effect is weakening in real time, as has been reported before (see, e.g., Wolk et al., 2013). Second, observe that PUM_THEMATICITY interacts with real TIME. We had said before that the favoring main effect of PUM_THEMATICITY is slightly puzzling. What we see now—consider the negative sign of the coefficient associated with the interaction term—is that the predictor's effect changes sign sometime after AD 1900 (as before, we simply add the value of the interaction coefficient to the coefficient of the main effect for each passing century). In Present-Day English, therefore, PUM_THEMATICITY does have the expected effect direction. Third, the negative coefficient associated with the interaction term FINAL_SIBILANCY = "yes" * TIME indicates that the disfavoring effect of final sibilancy on *s*-genitive choice is in fact growing stronger in real time, a finding that once again is compatible with the literature (Hinrichs & Szmrecsanyi, 2007).

In summary, given the huge literature on variation between the *s*-genitive and the *of*-genitive, Model 1 does not bear many surprises, some minor quirks notwithstanding.

Modeling the choice between the NN-genitive and the of-genitive (Model 2)

We next turn our attention to the first patch of uncharted territory in this study—the competition between the NN-genitive, as in (6), and the *of*-genitive, as in (7).

(6) **the NN-genitive**:
"He may start the day a little early," an [***embassy***]$_{\text{possessor}}$ [***spokesman***]$_{\text{possessum}}$ *said.* (1989tim1.n8b)

(7) **the *of*-genitive**:
This took on a new lease of life on Tuesday after art students erected a 30 ft. high model of the Statue of Liberty facing a [***portrait***]$_{\text{possessum}}$ **of** [***Mao Tse-tung***]$_{\text{possessor}}$ across the square. (1989tim1.n8b)

One would expect that this alternation exhibits some parallels with the variability between the *s*-genitive and the *of*-genitive (Model 1), thanks to the fact that the NN-genitive, like the *s*-genitive, places the possessor first (again, we remind the reader that we use the possessor/possessum terminology without making claims about the semantic nature of genitive relations). The difference is that possessor animacy—which is a factor historically associated with the *s*-genitive—is not expected to matter in Model 2, nor should final sibilancy matter as there is no haplology effect to be had with the NN-genitive.

Just like Model 1, Model 2 restricts attention to NN-genitives and *of*-genitives that are interchangeable with each other but not with the third option, the *s*-genitive.[11] The dataset on which Model 2 draws covers $n = 2832$ genitives ($n_{\text{NN-genitive}} = 905$, $n_{of\text{-genitive}} = 1927$). Both fitters independently found the exact same minimal adequate model. Model 2's index of concordance C is .96, indicating an excellent fit. The model correctly predicts 89% of all outcomes vis-à-vis a baseline prediction accuracy of 68%. There are no worries about multicollinearity (κ = 4.6), and the fixed-effect estimates discussed in the following are stable under bootstrap validation.

Table 3 reveals that the lexical by-item random effects in Model 2, pum_head_noun and, in particular, por_head_noun, are vastly more important than the random effects filename and register. As for register, science texts and letters receive a negative intercept adjustment; news texts receive a positive intercept adjustment.

Let us now discuss the fixed-effect structure of Model 2 (see Table 2). Increasing possessor length (por_length_words) disfavors usage of the NN-genitive, in accordance with the principle of end weight. Surprisingly, however, increasing possessum length (pum_length_words) usage *also* disfavors usage of the NN-genitive—the principle of end weight would seem to predict a favoring effect, as the NN-genitive places the possessum in second position. As expected, possessor animacy is not a significant predictor; alpha-persistence-related predictors also did not make it into the minimal adequate model. There is one beta-persistence-related predictor that is part of the minimal adequate model, beta_persistence_NN = "yes": recent usage of an NN-sequence favors reuse of an NN-genitive, as expected. As in Model 1, possessor thematicity is not significant while possessum thematicity (pum_thematicity) is: more frequent possessums favor usage of the NN-genitive. This effect is unexpected: more frequent (and, thus, more easily retrievable) possessums ought to favor the *of*-genitive as the option that places the possessum first (Easy First, MacDonald, 2013). We are also surprised to see that final sibilancy in the possessor (final_sibilancy) significantly discourages usage of the NN-genitive. This is unexpected because unlike the *s*-genitive, the NN-genitive—which has no overt genitive marker—cannot trigger a haplology effect. Time has the expected effect: as real time goes by, the odds for choice of the NN-genitive increase. There is one significant interaction term between time and a language-internal predictor, final_sibilancy, which indicates that the unexpected disfavoring effect of final sibilancy on NN-genitive-choice is weakening a bit in real time. That said, it does not go away.

As an interim summary, we conclude that Model 2 does in some ways exhibit the traits one would expect it to have. We note here, for example, the disfavoring effect of long possessors on NN-genitive choice, and the irrelevance of animacy. In other respects, Model 2 bears surprises—chief among them the fact that long possessums and possessors ending in a final sibilant repel the NN-genitive. In terms of language change, Model 2 shows that the NN-genitive is becoming more popular in real time, regardless of the weighting of the predictors.

Modeling the choice between the NN-genitive and the s-genitive (Model 3)

This section deals with variation between the NN-genitive, as in (8), and the *s*-genitive, as in (9).

(8) **the NN-genitive**:
For the [***vessel***]$_{\text{possessor}}$ [***walls***]$_{\text{possessum}}$, *a standard area was swabbed with sterile cotton wool . . .* (1975gibb.s8b)

(9) **the *s*-genitive**:
The union led by Mr Moss Evans must take the main responsibility for wrecking [***Labour***]$_{\text{possessor}}$**'s** [***plans***]$_{\text{possessum}}$ *to keep inflation down to single figures.* (1979stm1.n8b)

As for predictions, one would not expect to see any length effects in this particular alternation, as both variants establish possessor-before-possessum order. However, according to the literature (e.g., Rosenbach, 2006:90), animate possessors should disfavor the NN-genitive.

Model 3 restricts attention to NN-genitives and *s*-genitives that are interchangeable with each other, but it also includes—by contrast to Model 1 and Model 2—NN-genitives and *s*-genitives that are interchangeable with the third variant, the *of*-genitive. The dataset on which Model 3 will be based is therefore defined less restrictively than the other datasets. The reason is that applying the same standards as in the previous datasets would have led to crippling data sparsity.[12] Model 3's dataset consists of 676 genitives ($n_{\text{NN-genitive}} = 563$, $n_{s\text{-genitive}} = 113$).[13]

In finding the minimal adequate model, the two fitters encountered one disagreement concerning the main effect of FINAL_SIBILANCY and its interaction with TIME: one of the fitters included both the main effects and the interaction, the other fitter none. The two fitters agreed after discussion that both the (insignificant) main effect and the interaction term be included in the minimal adequate model. Model 3's index of concordance *C* is an excellent .94. Model 3 correctly predicts 91% of all outcomes vis-à-vis a baseline prediction accuracy of 83%. Multicollinearity is negligible ($\kappa = 2.42$). Except for the main effect of FINAL_SIBILANCY (which is not significant in the regression analysis anyway), the fixed-effect estimates discussed are stable under bootstrap validation. As far as the random effect structure is concerned, the by-item lexical factor POR_HEAD_NOUN and the by-subject effect FILENAME are comparatively important, while PUM_HEAD_NOUN and REGISTER do not have substantial effects.

We move on to a discussion of the fixed-effects structure (see Table 2). Contrary to expectations, we find that increasing possessum length (PUM_LENGTH_WORDS) significantly disfavors the NN-genitive, exactly as in Model 2 (recall that as we are in this section not dealing with a word order alternation, length effects are unexpected). But as predicted, animate possessors disfavor the NN-genitive and favor the *s*-genitive. Real TIME once again has a significant main effect indicating that the NN-genitive has been spreading at the expense of the *s*-genitive. There are moreover two significant interaction terms involving real TIME. First, thematicity of the possessum (PUM_THEMATICITY) has come to increasingly favor the NN-genitive vis-à-vis the *s*-genitive. Second, as in Model 2, final sibilancy in the possessor (FINAL_SIBILANCY) disfavors the NN-genitive, though in Model 3 the main effect is not significant. However, also as in Model 2, this disfavoring effect is weakening in real time.

To sum up Model 3, we have seen that possessor animacy has the expected effect: animate possessors repel the NN-genitive and attract the *s*-genitive. Just as Model 2, Model 3 shows that the NN-genitive is becoming more popular in real time, regardless of the weighting of the predictors. Contrary to expectations, however (and again just as in Model 2), we find a constituent length effect indicating that longer possessums disfavor the NN-genitive. The persistence/priming factors are across the board not significant, thanks maybe to the fact that there are no distinct word order patterns that could serve as carriers of priming effects.

Modeling ternary variability (Model 4)

In this section, we explore those *s*-genitives, as in (10), *of*-genitives, as in (11), and NN-genitives, as in (12), that are in competition with *both* other variants.

(10) **the *s*-genitive**:
The inquiry is the first big test for one of the [***CIA***]$_{\text{possessor}}$**'s** [***innovations***]$_{\text{possessum}}$ . . . (1989tim2.n8b)

(11) **the *of*-genitive**:
[***Beatings***]$_{\text{possessum}}$ **of** [***university students***]$_{\text{possessor}}$ *have been reported in the same region at Asyut and at Minya* . . . (1979stm1.n8b)

(12) **the NN-genitive**:
[***Union***]$_{\text{possessor}}$ [***negotiators***]$_{\text{possessum}}$ *hoped for a breakthrough at Hull on Friday* . . . (1979obs1.n8b)

The dataset consists of 2927 doubly interchangeable genitives ($n_{\text{NN-genitive}} = 470$, $n_{\textit{of}\text{-genitive}} = 2351$, $n_{\textit{s}\text{-genitive}} = 106$). To model the variation, we used "classical" binary logistic regression modeling with mixed effects to predict the odds for usage of the NN-genitive, as in (12), vis-à-vis its competitors in (10) and (11).[14]

Both fitters arrived independently at the exact same minimal adequate model. *C* is .96, which indicates excellent fit. The model accurately predicts 93% of all outcomes (baseline: 84%). There are two interaction effects in the minimal adequate model that did not survive bootstrap validation, namely PUM_LENGTH_WORDS * TIME, and PUM_THEMATICITY * TIME; all other effects discussed are stable. The lexical by-item random effects in Model 4—PUM HEAD_NOUN and, in particular, POR_HEAD_NOUN—are considerably more important than the random effects FILENAME and REGISTER. As for REGISTER, science texts receive a negative intercept adjustment, while news and, in particular, letters receive a positive intercept adjustment.

There are three fixed effects in Model 4 (see Table 2) that have the theoretically expected effects. First, longer possessors (POR_LENGTH_WORDS) disfavor the NN-genitive. This is compatible with the principle of end weight, in that the non-NN-genitives mostly consist of *of*-genitives, which place the possessor last. Second, recent usage of a generic NN-sequence (BETA_PERSISTENCE_NN = "yes") favors reuse of NN-genitives. Third, the NN-genitive is diachronically on the rise, as can be seen from the positive coefficient associated with the main effect of TIME.

The remainder of Model 4's fixed-effects structure has a number of interesting twists. Consider PUM_LENGTH_WORDS: precisely as in Models 2 and 3, increasing possessum length disfavors the NN-genitive (this is surprising, given that in Model 4 the non-NN-genitives mostly consist of *of*-genitives, which place the possessum first—it is these that should disfavor the NN-genitive). In addition, the effect's interaction with TIME appears to indicate that the disfavoring effect of PUM_LENGTH_WORDS strengthens diachronically. The significant, positive effect of ALPHA_PERSISTENCE_OF indicates that recent usage of an interchangeable *of*-genitive increases the odds for an NN-genitive, which is a bit curious. At the same time, we observe that while the main effect of ALPHA_PERSISTENCE_S is not significant, the interaction term involving TIME shows that a preceding *s*-genitive increasingly favors NN (probably because the *s*-genitive exhibits the same constituent order as the NN-genitive). Moving on, we note that Model 4 is the only model in this study where possessor thematicity (POR_THEMATICITY) has a significant effect (as a main effect and in interaction with TIME). The sign of the main effect is unexpected, in that highly thematic possessors disfavor the NN-genitive (given that the alternative to an NN-genitive is more often than not an *of*-genitive, possessor thematicity should favor the NN-genitive, thanks to the Easy First principle in MacDonald [2013]). That said, the interaction with TIME suggests that the possessor thematicity effect changes sign at around AD 1900, so that after AD 1900 possessor thematicity has the theoretically expected effect. However, increased possessum thematicity (PUM_THEMATICITY) favors the NN-genitive, whereas it was expected to disfavor. Finally, as in the other models involving the NN-genitive, final sibilancy in the possessor (FINAL_SIBILANCY) discourages usage of the NN-genitive.

In all, Model 4 offers some interesting clues as to how language users choose between the NN-genitive and its competitors. The NN-genitive is favored when a generic NN-sequence had been used recently, when the possessum is thematic, and in more modern texts. The NN-genitive is disfavored with long possessors and possessums, and when the possessor ends in a final sibilant. The effects of possessor thematicity and some persistence-related effects are diachronically unstable.

Possessor animacy is not a significant predictor in Model 4, though we hasten to add that there are few animate possessors in the dataset (200 of 2927). Of these 200 animate possessors, a full 42 (i.e., 21%) are encoded by the *s*-genitive, though its overall share is less than 4%. Thus there are corners in the datasets where animacy clearly plays a role: animate possessors attract the *s*-genitive.

DISCUSSION

Statistically speaking, the regression models are all quite good ones, with *C* values at or above .94 and with the models predicting 89% or higher of linguistic outcomes correctly (against varying baselines). Given that we derived our predictors from the literature on *s*-genitive versus *of*-genitive variation and that we did not consider some well-known but hard to annotate factors such as semantic relation, it is indeed noteworthy that the models predicting NN-genitive usage (Models 2–4) work so well.

But beyond statistical fit, what are this study's key substantive insights about genitive variation? First, we summarize what we have learned about the four variable contexts. We then review the extent to which the predictors we have studied have the theoretically expected effects.

The variable-context centered perspective: Profiling top predictor sets

To highlight differences between the four variable contexts analyzed in this paper, we adopt a more synoptic perspective: which predictors are important in particular variable contexts, but not in others? To address this question, we examine the regression models presented earlier from a different angle. Rather than focusing on regression coefficients, we rank the predictors according to their relative explanatory importance, using the chi-squared test statistics, which are output by likelihood ratio tests. These tests compare full minimal adequate models with all significant predictors (as reported in Table 2) to models with all predictors minus the predictor whose importance is to be established (see e.g., Speelman, Heylen, & Grondelaers, forthcoming).

Table 4 presents the top 3 explanatory predictors in each variable context. We begin by discussing the variation between the s-genitive and the *of*-genitive (Model 1). Given the extensive literature on this variation, the ranking in Table 4 bears no major surprises. Possessor animacy is the most important predictor by far, followed by possessor length and final sibilancy. In the variation between the NN-genitive and the of-genitive (Model 2), it is, by contrast, possessum length that is identified as the most important predictor. The main effect of real time is the second most important predictor in this variable context, which indicates robust diachronic change. In variation between the NN-genitive and the *s*-genitive (Model 3), possessor animacy is the top-ranked predictor (recall that the *s*-genitive is preferred when the possessor is animate), followed by the main effect of real time, indicating once again that this variable context is subject to massive change. Given that we are not dealing with a word order alternation, it is once again surprising that possessum length is relatively crucial. As for variation between the NN-genitive and non-NN-genitives (either the *of*-genitive or the *s*-genitive), which was analyzed in Model 4, possessor length comes out on top (remember that long possessors, like long possessums, disfavor the NN-genitive), followed by final sibilancy in the possessor (likewise disfavoring the NN-genitive) and possessum thematicity (which favors the NN-genitive). Note that the main effect of real time is relatively important (rank 4, not displayed in the table) and is not far behind the three top-ranked predictors.

In summary, animacy takes center stage in Models 1 and 3, which both involve the *s*-genitive. Length effects are important throughout, while real time is strongly implicated in those models where the NN-genitive is involved.

TABLE 4. *Relative importance of predictors in Models 1–4 according to likelihood ratio tests*

	Predictor	Model Chi-square
s-genitive vs. *of*-genitive (Model 1)	1. ANIMACY	371.6
	2. POR_LENGTH_WORDS	111.0
	3. FINAL_SIBILANCY	70.8
of-genitive vs. NN-genitive (Model 2)	1. PUM_LENGTH_WORDS	126.4
	2. TIME	92.8
	3. FINAL_SIBILANCY	68.0
NN-genitive vs. *s*-genitive (Model 3)	1. ANIMACY	45.6
	2. TIME	15.9
	3. PUM_LENGTH_WORDS	7.8
NN-genitive vs. non-NN-genitives (Model 4)	1. POR_LENGTH_WORDS	42.7
	2. FINAL_SIBILANCY	38.6
	3. PUM_THEMATICITY	31.6

Note: Displayed: top three predictors. POR = possessor, PUM = possessum.

The predictor-centered perspective: Significance and effect directions

An alternative perspective is to synthesize what we have learned about the *predictors* that predict variant choice in these contexts (rather than the variable contexts per se). We begin by discussing constituent length (POR_LENGTH_WORDS and PUM_LENGTH_WORDS). Possessum length is a significant predictor throughout; possessor length is not significant in Model 3. A key finding is that increased constituent length, when significant, consistently disfavors the NN-genitive, in contrast to the predictions of previous research on weight effects. That is, the NN-genitive is preferred when both the possessor and possessum constituents are rather light. Thus, across the four contexts, mean length of the possessor plus possessum phrase (in words) is 3.3 words for the *of*-genitive, 2.9 words for the *s*-genitive, but only 2.4 words for the NN-genitive. If we reasonably assume that heavier constituents reflect cognitive complexity, and that the NN-genitive—because it does not have an overt genitive marker and hence lacks a structural signal of the boundary between the possessor and possessum—is less explicit than the other genitive variants, these differentials are predicted by Rohdenburg's (1996) complexity principle: explicit options (*s*-genitive, *of*-genitive) are preferred in more complex environments.

Possessor animacy (ANIMACY) works overall as advertised. It has the expected effect in Model 1 (where animate possessors favor the *s*-genitive; see Rosenbach [2005]) and in Model 3 (where animate possessors favor the *s*-genitive; see Rosenbach [2006:90, 2007a:177, 2007b:166]. It is not a significant predictor in Model 2 or in Model 4, although we have seen that in the ternary dataset, animate possessors do seem to attract the *s*-genitive.

For reasons that are currently unclear, the three alpha-persistence predictors (ALPHA_PERSISTENCE_OF, ALPHA_PERSISTENCE_S, ALPHA_PERSISTENCE_NN) do not constrain variation in the expected fashion. In Models 2 and 3, none of the predictors is selected as significant; in Models 1 and 4, the factors do have significant but counterintuitive effects. By contrast, the beta-persistence predictors have on the whole the theoretically expected effects: recent usage of similar material primes genitive variants.

We included type-token ratio (TTR) in our predictor portfolio because Hinrichs and Szmrecsanyi (2007) found that higher type-token ratios predict choice of the *s*-genitive in 20th-century press English. We did not find a corresponding effect in any of the regression models.

As for thematicity, remember that we measured the local text frequency of the possessor (POR_THEMATICITY) *and* the possessum head (PUM_THEMATICITY) in the corpus text under analysis (following Hinrichs & Szmrecsanyi, 2007). We predicted that more frequent, and thus accessible, "easy" constituents should attract a genitive variant that places the accessible constituent first (Easy First, MacDonald, 2013). In the event, it turned out that possess*or* thematicity is not a significant predictor in Models 1–3, while in Model 4, where it is significant, it does not have the theoretically expected effect direction. So possessor thematicity does not play a huge role in our data. Possess*um* thematicity seems to be overall more influential: the pattern is that in models where the NN-genitive is involved, increased possessum thematicity favors the NN-genitive. This does seem to go against the Easy First principle (MacDonald, 2013), but then again, the effect may be in line with Rohdenburg's Complexity Principle: if the possessum is thematic (textually frequent), it is arguably accessible and thus cognitively less complex; under such circumstances, language users may feel safe in using the less explicit NN-genitive. We note that the Complexity Principle–based account would also explain the favoring effect of possessum thematicity on *s*-genitive choice in Model 1, insofar as the *s*-genitive is the less explicit coding option vis-à-vis the more verbose *of*-genitive construction: thematic possessums can be seen as noncomplex, hence they take the less explicit genitive variant. In summary, the possessum thematicity effect is robust and constitutes an important characteristic of the NN-genitive.

Final sibilancy in the possessor had the expected effect in Model 1, where it discourages *s*-genitive usage. In Models 2 and 4, however, final sibilancy disfavors the NN-genitive. This is at first view a bit mysterious, because there is no haplology effect to be had with the NN-genitive. Notice though that in coding for final sibilancy, we did not distinguish between plural markers and final sibilants in the root of a possessor head noun. Thus, the final sibilancy effect is probably a plurality effect—premodifier nouns do not normally have plural inflection (Feist, 2012:269).

We modeled two lexical effects, POR_HEAD_NOUN and PUM_HEAD_NOUN, as random effects (adjustments to the intercept). These are robustly implicated in genitive variation throughout. In Models 1–3, the possessor head noun lexemes are more important attractors than the possessum head noun lexemes. In Model 4, possessum head noun lexemes are more important attractors.

We finally turn to a discussion of the two language-external factors that we considered. Real time (TIME) we modeled as a fixed effect that was allowed to interact with language-internal predictors, which it did extensively. As for the main effect of real time, the conclusion is that TIME is a significant predictor in all four models: in models where the NN-genitive is involved, it is diachronically on the rise.

By-subject (i.e., by-author) variation, which we operationalized by modeling FILENAME as a random effect, is typically less important than lexical effects but more important than register differences. The variable REGISTER distinguished science texts, news texts, and letters via a random effect. In Models 2 and 4, science texts received negative intercept adjustments, while the news register received a positive intercept adjustment. This can be interpreted to mean that in a variationist perspective, science writers are less inclined, all other things being equal (!), to use the NN-genitive than writers in other registers are. Yet, it remains true that the NN-genitive is frequent in absolute terms in the science register (see also, e.g., Biber & Gray, 2016)—but crucially we also find lots of tokens in science texts where writers could have opted for an NN-genitive but did not. So here we are dealing with an aspect of genitive variability where the variationist method (which is concerned with variant rates and linguistic choice-making) and the text-linguistic method (with its focus on the overall frequency of features in texts) put in relief different facets of the same reality.

CONCLUDING REMARKS

Using a large, richly annotated dataset covering three written registers throughout the LateModE period—the time slice that we knew saw the evolution of the NN-genitive—our study had two major objectives: first, to investigate how language users choose between the NN-genitive and competing genitive variants; second, to showcase ways to handle a complex dataset about nonbinary grammatical variation with a number of partially overlapping variable contexts. As for the first objective, key findings include that there is an overall drift toward the NN-genitive, and that the NN-genitive is popular with short constituents, inanimate possessors, highly thematic possessums, and nonplural possessors. We also saw that there are robust probabilistic changes during the LateModE period. On the methodological plane, we have addressed the complexities of the dataset by splitting it up into four subdatasets, each one covering a different variable context that can be analyzed using regression analysis.

As always, there are numerous ways in which this research can be extended. For example, it would be desirable to use a more fine-grained animacy scheme. With regard to final sibilancy, future research should distinguish between root *s* and plural *s*. There are also some potentially important predictors that we did not consider at all, such as definiteness/referentiality, nature of the genitive relation, information status (given versus new), and rhythmic well-formedness (Shih, Grafmiller, Futrell, & Bresnan, 2015). And finally, future study should explore in more detail thematicity effects, including the explanatory power of relative measures of possessor/possessum thematicity.

Let us conclude with some philosophical afterthoughts: What is a genitive in English? How should we define the scope of a linguistic variable? And what is the appropriate domain of inquiry for studies of linguistic variation? Traditionally, the domain of inquiry for studies of genitive variation has been restricted to those linguistic tokens that are interchangeable between the *s*-genitive and the *of*-genitive. The linguistic status of noninterchangeable tokens is not clear; they have simply been disregarded as not relevant to the domain of inquiry. But this is another way of saying that at present, we know little about another type of variability: why some tokens are interchangeable and others are not—which is yet another issue that future research should address.

Restricting in this way the domain of inquiry raises methodological issues when we consider variables with three variants: is there a basic variant (the *s*-genitive, say) that should be required as part of the definition of the linguistic variable, or should we include all tokens that are interchangeable with any two of the three possible variants? Notice that our analysis has included tokens that would have been excluded in previous analyses: occurrences of the *s*-genitive that are interchangeable with the NN-genitive but *not* with the *of*-genitive, and occurrences of the *of*-genitive that are interchangeable with the NN-genitive but *not* with the *s*-genitive (see Table 1 for examples). We have further included occurrences of the NN-genitive that are interchangeable with only one of the other two variants. It is not clear that alternative operational definitions would have been preferable. For example, it is not feasible to operationally define the genitive on semantic grounds (very few actual genitive tokens express the classical meaning of "possession"). Furthermore, there seems to be no good reason for privileging one of the genitive variants as the "basic" one and requiring interchangeability with that variant as part of the operational definition of "genitive." Rather, what we seem to be dealing with is a web of variation, with different degrees of interchangeability among the variants that make up that web.

However, once the lid is cracked on that Pandora's box, we are forced into a very different perspective on the nature of linguistic variation. For example, in the case of genitive variation, we find interchangeability with variants that have not been considered here. In some cases, those new variants are clearly interchangeable with all of the other traditional genitive variants, as in (13):

(13) *FBI director—FBI's director—director of the FBI—director from the FBI*

But many other instances are interchangeable with only one or two of the other variants. Consider (14) to (16):

(14) *Energy Department spokesman—Energy Department's spokesman—spokesman of the Energy Department* ??—*spokesman for the Energy Department—spokesman from the Energy department*

(15) *federal funding restriction—restriction on federal funding—restriction of federal funding—federal funding's restriction* ??

(16) *oil price increase—increase in oil prices—oil price's increase* ??—*increase of oil prices* ??

Research is underway to explore this larger web of variation. We are also exploring the linguistic characteristics of structures that are *not* interchangeable with other variants. That is, rather than simply discarding these tokens as theoretically uninteresting, we are trying to determine the linguistic characteristics of noninterchangeable structures, and describe how those characteristics compare/contrast to interchangeable variants. Our ultimate goal is to reconcile the findings from these different research perspectives, accounting for the full range of phenomena associated with a domain of linguistic variation.

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit http://dx.doi.org/10.1017/S0954394515000198.

NOTES

1. All examples in this paper are drawn from *A Representative Corpus of Historical English Registers* (ARCHER; Biber, Finegan, & Atkinson, 1994; Yáñez-Bouza, 2011) and are referenced by ARCHER text identifiers. Note that we use the customary labels "possessor" and "possessum" to refer to the constituents of genitive constructions without wishing to make any claims about the exact semantic relationship between the constituents.

2. We follow Feist (2012:293) in considering the NN-genitive a true genitive. What we call "NN-genitives" are also known as "noun + noun sequences" (Rosenbach, 2006), or "noun + noun constructions" (Rosenbach, 2007a) in the literature.

3. In this study, we use "predictor" as a cover term referring to both categorical conditioning factors (such as possessor animacy) and to continuous variables (such as constituent length in words).

4. Instead of relying on two coders and measuring intercoder reliability, an alternative approach to coding genitives for interchangeability would be an acceptability judgment study with a larger number of participants (e.g., $n = 20$).

5. We chose to model the explanatory variable REGISTER as a random effect and not as a fixed effect because in this particular study we are not explicitly interested in the statistical significance of register differences and/or in interaction effects between REGISTER and other predictors. Under such circumstances, regardless of the conceptual status of the factor, the analyst may be justified in modeling a factor as a random effect in order not to waste precious degrees of freedom (Zuur et al., 2009:106). Nonetheless, we experimented with modeling REGISTER as a fixed effect. In the model for the *s*-genitive versus the *of*-genitive (Model 1), modeling REGISTER as a fixed effect results in the interaction between possessum thematicity (PUM_THEMATICITY) and TIME losing significance. Apart from this, the regression models with REGISTER as a random effect are almost identical to the models with REGISTER as a fixed effect.

6. We use R version 3.0.1 and lme4 version 0.999999-2 (R Development Core Team, 2013).

7. Outside linguistics, see, for example, Burnham and Andersen (2002) for a view that considers model selection an integral part of the data analysis endeavor.

8. To remove collinearity, continuous variables were centered around their mean using the c.() function (see http://hlplab.wordpress.com/2011/02/24/diagnosing-collinearity-in-lme4/).

9. That said, note that we experimented with a model that includes genitives that are in principle also interchangeable with the NN-genitive ($n = 6652$). This model has the exact same set of significant predictors, which have the exact same effect directions.

10. Here and throughout the paper, κ was measured using the function collin.fnc() in R package languageR (version 1.4).

11. We again experimented with a model that includes genitives that are in principle also interchangeable with the *s*-genitive ($n = 5653$). The predictors that are significant in the more restricted model are also significant in the less restricted model, and in addition, the signs are identical throughout.

12. Restricting attention to NN-genitives and *s*-genitives that are interchangeable with each other but not with the *of*-genitive yields a small and extremely lopsided dataset ($n = 100$, $n_{\text{NN-genitive}} = 93$, $n_{s\text{-genitive}} = 7$) that is not amenable to quantitative analysis.

13. It has been suggested that the maximal number of parameters (i.e., coefficients) in a logistic regression model should not exceed the frequency of the less frequent outcome divided by 10 (Hosmer & Lemeshow, 2000:346–347). According to this criterion, Model 3 is fine (7 parameters, $n = 676$ observations, $n[s\text{-genitive}] = 113$).

14. We experimented with multinomial and polytomous regression techniques (as implemented in the R libraries nnet and polytomous) that can in principle be used to model nonbinary linguistic outcomes, but found that the resulting models were inadequate because *s*-genitives, which are relatively rare in the dataset, were consistently misclassified.

REFERENCES

Baayen, R. Harald. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press.

Bates, Douglas, Maechler, Martin, & Bolker, Ben. (2012). *lme4: Linear mixed-effects models using S4 classes*. Available at: http://cran.r-project.org/web/packages/lme4/index.html. Accessed December 10, 2015.

Baugh, John. (1980). A re-examination of the Black English Copula. In W. Labov (ed.), *Locating language in time and space*. New York: Academic Press. 83–106.

Behaghel, Otto. (1909). Beziehungen zwischen Umfang und Reihenfolge von Satzgliedern. *Indogermanische Forschungen* 25:110–142.

Biber, Douglas, Finegan, Edward, & Atkinson, Dwight. (1994). ARCHER and its challenges: Compiling and exploring A Representative Corpus of Historical English Registers. In G. Fries, U. Tottie, & P. Schneider (eds.), *Creating and using English language corpora: Papers from the Fourteenth International Conference on English Language Research and Computerized Corpora*. Amsterdam: Rodopi. 1–13.

Biber, Douglas, & Gray, Bethany. (2011). Grammatical change in the noun phrase: The influence of written language use. *English Language and Linguistics* 15(2):223–250.

______. (2013). Being specific about historical change: The influence of sub-register. *Journal of English Linguistics* 41(2):104–134.

______. (2016). *Grammatical complexity in academic English: Linguistic change in writing*. Cambridge: Cambridge University Press.

Biber, Douglas, Johansson, Stig, Leech, Geoffrey, Conrad, Susan, & Finegan, Edward. (1999). *Longman grammar of spoken and written English*. Harlow: Longman.

Burnham, Kenneth P., & Anderson, David Raymond. (2002). *Model selection and multimodel inference a practical information-theoretic approach*. 2nd ed. New York: Springer.

Ehret, Katharina, Wolk, Christoph, & Szmrecsanyi, Benedikt. (2014). Quirky quadratures: On rhythm and weight as constraints on genitive variation in an unconventional data set. *English Language and Linguistics* 18(2):263–303.

Feist, Jim. (2012). What controls the "genitive variation" in Present-Day English? *Studies in Language* 36(2):261–299.

Gelman, Andrew, & Hill, Jennifer. (2007). *Data analysis using regression and multilevel/hierarchical models*. Analytical Methods for Social Research. Cambridge: Cambridge University Press.

Grafmiller, Jason. (2014). Variation in English genitives across modality and genres. *English Language and Linguistics* 18(3):471–496.

Gries, Stefan Th. (2005). Syntactic priming: A corpus-based approach. *Journal of Psycholinguistic Research* 34(4):365–399.

______. (2013). *Statistics for linguistics with R: A practical introduction*. 2nd rev. ed. Berlin: De Gruyter Mouton.

Hinrichs, Lars, & Szmrecsanyi, Benedikt. (2007). Recent changes in the function and frequency of Standard English genitive constructions: A multivariate analysis of tagged corpora. *English Language and Linguistics* 11(3):437–474.

Hosmer, David W., & Lemeshow, Stanley. (2000). *Applied logistic regression*. New York: Wiley.

Jankowski, Bridget L., & Tagliamonte, Sali A. (2014). On the genitive's trail: Data and method from a sociolinguistic perspective. *English Language and Linguistics* 18(2):305–329.

Labov, William. (1969). Contraction, deletion and inherent variability of the English copula. *Language* 45(4):715–762.

———. (1972). *Sociolinguistic patterns*. Philadelphia: University of Philadelphia Press.

———. (1982). Building on empirical foundations. In W. Lehmann & Y. Malkiel (eds.), *Perspectives on historical linguistics*. Amsterdam: Benjamins. 17–92.

MacDonald, Maryellen C. (2013). How language production shapes language form and comprehension. *Frontiers in Psychology* 4:226.

Meyerhoff, Miriam, & Walker, James A. (2007). The persistence of variation in individual grammars: Copula absence in "urban soujourners" and their stay-at-home peers, Bequia (St Vincent and the Grenadines). *Journal of Sociolinguistics* 11(3):346–366.

Osselton, Noel. (1988). Thematic genitives. In G. Nixon & J. Honey (eds.), *An historic tongue: Studies in English linguistics in memory of Barbara Strang*. London: Routledge. 138–144.

Pinheiro, José C., & Bates, Douglas M. (2000). *Mixed-effects models in S and S-PLUS*. New York: Springer.

Poplack, Shana, & Sankoff, David. (1987). The Philadelphia story in the Spanish Caribbean. *American Speech* 62(4):291–314.

Raumolin-Brunberg, Helena. (2005). The diffusion of subject you: A case study in historical sociolinguistics. *Language Variation and Change* 17:55–73.

R Development Core Team. (2013). *R: A language and environment for statistical computing*. Vienna, Austria. Available at: http://www.R-project.org/. Accessed December 10, 2015.

Rickford, John R., Ball, Arnetha, Blake, Renee, Jackson, Raina, & Martin, Nomi. (1991). Rappin on the copula coffin: Theoretical and methodological issues in the analysis of copula variation in African-American Vernacular English. *Language Variation and Change* 3:103–132.

Rohdenburg, Günter. (1996). Cognitive complexity and increased grammatical explicitness in English. *Cognitive Linguistics* 7(2):149–182.

Rosenbach, Anette. (2005). Animacy versus weight as determinants of grammatical variation in English. *Language* 81(3):613–644.

———. (2006). Descriptive genitives in English: A case study on constructional gradience. *English Language and Linguistics* 10(1):77–118.

———. (2007a). Exploring constructions on the web: A case study. In M. Hundt, N. Nesselhauf, & C. Biewer (eds.), *Corpus linguistics and the Web*. Amsterdam: Rodopi. 167–190.

———. (2007b). Emerging variation: Determiner genitives and noun modifiers in English. *English Language and Linguistics* 11(1):143–189.

———. (2014). English genitive variation—The state of the art. *English Language and Linguistics* 18 (2):215–262.

Sankoff, David, & Labov, William. (1979). On the use of variable rules. *Language in Society* 8:189–222.

Sankoff, David, & Rousseau, Pascale. (1989). Statistical evidence for rule ordering. *Language Variation and Change* 1(1):1–18.

Shih, Stephanie, Grafmiller, Jason, Futrell, Richard, & Bresnan, Joan. (2015). Rhythm's role in genitive construction choice in spoken English. In R. Vogel & R. Vijver (eds.), *Rhythm in cognition and grammar*. Berlin: De Gruyter. 207–234.

Speelman, Dirk, Heylen, Kris, & Grondelaers, Stefan. (forthcoming). A bottom-up, data-driven operationalization of semantic classes and predictability in syntactic alternation research. In S. Grondelaers & R. van Hout (eds.), *New ways of analyzing syntactic variation*. Berlin: De Gruyter Mouton.

Szmrecsanyi, Benedikt. (2006). *Morphosyntactic persistence in spoken English: A corpus study at the intersection of variationist sociolinguistics, psycholinguistics, and discourse analysis*. Berlin: Mouton de Gruyter.

Tagliamonte, Sali, & Baayen, R. Harald. (2012). Models, forests, and trees of York English: Was/were variation as a case study for statistical practice. *Language Variation and Change* 24 (2):135–178.

Travis, Catherine E. (2007). Genre effects on subject expression in Spanish: Priming in narrative and conversation. *Language Variation and Change* 19(2):101–135.

Wasow, Thomas. (2002). *Postverbal behavior*. Stanford: CSLI Publications.

Wolfram, Walt. (1974). The relationship of White Southern Speech to Vernacular Black English. *Language* 50(3):498–527.

Wolk, Christoph, Bresnan, Joan, Rosenbach, Anette, & Szmrecsanyi, Benedikt. (2013). Dative and genitive variability in Late Modern English: Exploring cross-constructional variation and change. *Diachronica* 30(3):382–419.

Yáñez-Bouza, Nuria. (2011). ARCHER past and present (1990–2010). *ICAME Journal* 35:205–236.

Zaenen, Annie, Carlette, Jean, Garretson, Gregory, Bresnan, Joan, Koontz-Garboden, Andrew, Nikitina, Tatiana, O'Connor, M. Catherine, & Wasow, Tom. (2004). Animacy encoding in English: Why and how. In D. Byron, & B. Webber (eds.), *Proceedings of the 2004 ACL workshop on discourse annotation, Barcelona, July 2004*. 118–125.

Zuur, Alain F., Ieno, Elena N., Walker, Neil J., Saveliev, Anatoly A., & Smith, Graham M. (ed.). (2009). *Mixed effects models and extensions in ecology with R.* Statistics for Biology and Health. New York: Springer.

Zwicky, Arnold M. (1987). Suppressing the Zs. *Journal of Linguistics* 23:133–148.