CHAPTER 6

The Role of Gender in Postcolonial Syntactic Choice-Making
Evidence from the Genitive Alternation in British and Sri Lankan English

Stefan Th. Gries, Benedikt Heller and Nina Funke

6.1 Introduction

6.1.1 General Introduction

There is a rich history of research on the English genitive alternation (Rosenbach 2014; also Chapter 9), the choice between the *s*-genitive and the *of*-genitive. The two options differ mainly in two respects: the ordering and linking of two constituents. In the *s*-genitive as in (1a), the possessor (*Sri Lanka*) precedes the possessum (*self-interest*), whereas it follows the possessum (*image*) in the *of*-genitive as in (1b). To indicate the genitive relation between the constituents, the clitic *'s* (or simply ' with plurals) is added to the possessor phrase in *s*-genitives, whereas in *of*-genitives, constituents are connected by the preposition *of*. Additionally, the definite article is not explicitly stated in *s*-genitives (since *'s* has a deterministic function), so if we wanted to transform (1a) into an *of*-genitive, we would have to add it (i.e. *the self-interest of Sri Lanka*).

(1) a. So let us now serve Sri Lanka$_{[\text{POSSESSOR}]}$'s self-interest$_{[\text{POSSESSUM}]}$ by making use of the potential that this whole region has for massive economic development in the years to come <ICE-SL: S1B-051#109:1:A>

b. You had been manufacturing films to tarnish the image$_{[\text{POSSESSUM}]}$ of Sri Lanka$_{[\text{POSSESSOR}]}$ or photographs with the aim of promoting communism, terrorism or creating communal tension . . . <ICE-SL:S1B-054#102:1:A>

The present investigation uses a variationist approach (e.g. Szmrecsanyi 2017) to grammatical choice-making and, in doing so, focuses on genitives that constitute 'alternate ways of saying "the same" thing' (Labov 1972: 188; for a discussion of sameness in the genitive alternation, see Rosenbach 2002).

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

This entails excluding certain forms, which cannot be expressed in the respective other form (Rosenbach 2002); these forms include appositive genitives (e.g. *the state of California*), descriptive genitives (e.g. *person of colour*), double genitives (e.g. *the friend of John's*), noun–noun genitives (e.g. *satellite photographs*; but see Szmrecsanyi et al. 2016), partitive genitives (e.g. *one of my friends*) and idiomatic genitives (e.g. *Valentine's Day*). In focusing on the so-called choice context (i.e. on cases that can be expressed in both genitive variants), we seek to produce results that complement recent multifactorial studies of genitive choice (e.g. Grafmiller 2014; Heller et al. 2017a).

Linguists have long been aware that the choice between the *s*-genitive and the *of*-genitive is subject to multiple constraints, the most important of which is possessor animacy. In most studies, this refers to the binary distinction between animate (2a) and inanimate (2b) possessors, but more fine-grained additions to this distinction have been found to be important in explaining genitive choice (e.g. Wolk et al. 2013): collective (2c), locative (2d) and temporal (2e) possessors. In essence, the higher a possessor is on the animacy scale, the more likely it is to be used in an *s*-genitive.

(2) a. What form of government will then replace Saddam Hussein$_{[\text{POSSESSOR}]}$'s dictatorship$_{[\text{POSSESSUM}]}$? <ICE-GB:W2E-001#57:3>
b. The boiling point$_{[\text{POSSESSUM}]}$ of coconut oil$_{[\text{POSSESSOR}]}$ is more conducive for frying <ICE-SL:S1A-006#154:1:A>
c. Had the timely arrival$_{[\text{POSSESSUM}]}$ of US and British forces$_{[\text{POSSESSOR}]}$ not prevented this manoeuvre, some 5 million barrels a day of oil production might have been put at risk <ICE-GB:W2E-001 #45:3>
d. Mahatma Gandhi called India Sri Lanka$_{[\text{POSSESSOR}]}$'s nearest neighbour$_{[\text{POSSESSUM}]}$ <ICE-SL:S1B-051#112:1:A>
e. Mani, now today$_{[\text{POSSESSOR}]}$'s question$_{[\text{POSSESSUM}]}$ again; let me repeat that . . .<ICE-SL:S1A-094#14:1:A>

Almost equally important is the length of the constituents (for a comparison of length and animacy, see Rosenbach 2005). If the possessor is particularly long (in relation to the possessum, that is), it is more likely to be used in an *of*-genitive, in which it follows the possessum (3a). On the other hand, if the possessum is particularly long, it is more likely to be used in an *s*-genitive, in which it, again, is placed in the final position (3b). In other words, genitive choice corresponds to the principle of end-weight (Behaghel 1909), according to which longer constituents are often placed last. Both the

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

effect of possessor animacy and the effect of constituent lengths reflect a more general principle of linguistic choice-making: Easy First (MacDonald 2013), which states that constituents that are more easily retrievable from memory (e.g. animate [see Bock 1982: 15 on egocentric bias] and/or short ones) tend to be placed first.

(3) a. Now Mr deputy speaker with the conclusion$_{[\text{POSSESSUM}]}$ of the conflict situation$_{[\text{POSSESSOR}]}$ <ICE-SL:S1B-056#26:1:A>
b. And also another interesting thing is Ernest Hemingway$_{[\text{POSSESSOR}]}$'s last posthumously published novel$_{[\text{POSSESSUM}]}$ <ICE-SL:S2A-025#72:1:A>

Beyond possessor animacy and constituent length, the present study considers several additional constraints: three language-internal (i.e. sibilancy, definiteness and semantic relation) and three language-external ones (i.e. modality, variety and gender). Sibilancy refers to the presence of a final sibilant at the end of the possessor phrase. If a sibilant ([s], [z], etc.) is present there, language users tend to avoid the *s*-genitive because the combination of the sibilant and the clitic *'s* creates a repetitive sound sequence (4). Definiteness also refers to the possessor and simply distinguishes definite and indefinite ones; if the possessor is definite (5a), *s*-genitive usage is more probable than with indefinite possessors (5b). Semantic relation is here operationalised as a binary distinction between prototypical (including part–whole, kinship and legal) relations (e.g. [6a]) and non-prototypical ones (6b). With prototypical semantic relations, the *s*-genitive is usually more frequent than with non-prototypical ones. The effects of these three language-internal constraints have been found and replicated in many studies (an overview of which can be found in the appendix to Rosenbach 2014).

(4) In Peter Lucius$_{[\text{POSSESSOR}]}$'s terms$_{[\text{POSSESSUM}]}$ ... <ICE-SL:S2B-049#29:1:A>
(5) a. we met auntie Ashani$_{[\text{POSSESSOR}]}$'s daughter$_{[\text{POSSESSUM}]}$ ... <ICE-SL:S1A-011#38:1:B>
b. six hundred years ago, the delicious flavour$_{[\text{POSSESSUM}]}$ of mushrooms$_{[\text{POSSESSOR}]}$ intrigued the Pharaohs of Egypt <ICE-SL:S1B-021#138:1:D>
(6) a. the slave owners$_{[\text{POSSESSOR}]}$' children$_{[\text{POSSESSUM}]}$ ... <ICE-SL:S2A-024#115:1:A>
b. you know the underlying principles$_{[\text{POSSESSUM}]}$ of all the grammar issues$_{[\text{POSSESSOR}]}$ <ICE-SL:S1A-015#35:1:A>

The language-external predictors that are included are modality, variety and gender. Modality (*spoken* vs. *written*) is included as a control because it has been found to have an effect, especially in interaction with language-internal predictors (e.g. Grafmiller 2014). Szmrecsanyi and Hinrichs (2008) report more *s*-genitive usage in spoken texts, which they attribute to a difference in average formality (with spoken language usually being less formal). Low formality, in turn, has long been recognised as a factor that increases the use of the *s*-genitive (Altenberg 1982). Genitive choice across different varieties has only recently been studied in detail (Heller et al. 2017a; Szmrecsanyi et al. 2017; Heller 2018), but results show significant differences, especially for variety as a moderator of possessor animacy. In essence, it was found that possessor animacy triggers the *s*-genitive more strongly in Inner-Circle varieties (e.g. Britain, Canada; Kachru 1985) than in Outer-Circle varieties (e.g. Heller et al. 2017b). Gender has – to our knowledge – so far not been studied as predictor of genitive choice. A study of the (arguably comparable) dative alternation that included gender found non-significant results (Kendall et al. 2011).

The present study seeks to add to the research of syntactic alternations across varieties by investigating genitive choice in Sri Lankan English (SriLE). SriLE is a postcolonial variety with a clear variety-specific structural profile that spans several linguistic levels (e.g. Meyler 2007; Künstler et al. 2009; Bernaisch 2012, 2015). Within Schneider's Dynamic Model, SriLE has thus passed the nativisation phase and is arguably on its way towards endonormative stabilisation (see Mukherjee 2008: 361). In order to characterise genitive choice in SriLE, we compare it to its historical input-variety British English (BrE).

Thus far, there has been little research on gender differences in SriLE. In his dictionary of SriLE, Meyler (2007: 53) presents words that women use more often, such as *child!* as 'a colloquial term of address'. Bernaisch (2012) investigates Sri Lankans' attitudes towards American, British, Indian and SriLE without finding a statistically significant difference between female and male participants. Revis and Bernaisch (in press) report a significant effect of gender on the choice of filled and unfilled pauses in a conditional inference tree. However, in their general linear mixed-effects model, gender did not reach significance. Gunesekera (2005), in her analysis of the postcolonial identity of SriLE, found the phenomenon of topicalisation to be more common in female than in male speech, while sports metaphors and swear words were – at least in public – limited to male speech (see Gunesekera 2005: 137).

With respect to genitive choice, SriLE has not yet received much scholarly attention (pace Heller et al. 2017a). Given that previous studies of genitive choice across varieties revealed that language-external factors might influence genitive choice only as moderator of language-internal constraints (e.g. variety, which was found to moderate the strength of possessor animacy) and given the rich history and importance of the study of gender-based differences in language use (see the introduction to this volume), we reckon that the systematic study of gender in genitive choice constitutes a research gap. The present study, therefore, seeks to complement the current body of research on the genitive alternation by 1) including gender as language-external variable and 2) by investigating gender across varieties.

6.1.2 Overview of the Present Chapter

In the following section on methods (Section 6.2), we will first outline the data extraction process and show some descriptive statistics of the above-mentioned predictors (Section 6.2.1). In Section 6.2.2, we will present the details of the statistical analysis, whose results are then described and visualised in Section 6.3. In the final section, Section 6.4, we will provide a discussion and concluding remarks.

6.2 Methods

6.2.1 Data

Data were extracted from a 10% register-stratified sample of the British component of the International Corpus of English (ICE; Greenbaum 1996) and a 25% sample of the Sri Lankan component (ICE-SL) because at the time of our retrieval only 25% of the spoken component was available. The extraction of interchangeable genitives was accomplished in several (partly automatised) steps: first, all text units containing genitive markers (i.e. of, 's and –s') were automatically extracted. Then, an automatic classification determined interchangeability. Annotation of predictors was done automatically where possible, based on computational work from Heller (2018). In every step, manual corrections were made where necessary. Speaker gender information for each case was taken from the respective metadata in the case of ICE-SL and from metadata made available by Martin Schweinberger in the case of ICE-Great Britain (ICE-GB).[1]

[1] Available at www.martinschweinberger.de/blog/resources/.

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

Table 6.1 *Composition of the data with respect to* VARIETY *and* GENDER

VARIETY	GENDER	GENITIVE: *of*	GENITIVE: *s*	Sum
Great Britain	*Male*	305	104	409
	Female	48	19	67
	Unknown	434	153	587
Sri Lanka	*Male*	1,185	298	1,483
	Female	580	310	890
	Unknown	500	109	609
Sum (baseline %)		3,052 (75.5%)	993 (24.5%)	4,045

Altogether, we ended up with 4,045 cases that are distributed across VARIETY and GENDER as shown in Table 6.1, with an overall preference for *of*- over *s*-genitives to a degree that raised initial concerns about the class imbalance problem (the problem that if the [two] levels of the dependent variables are very skewed in favour of one option already, it can become problematic to get good results out of a regression).

In addition to the variables shown in Table 6.1, the data were then annotated with regard to several other predictors discussed in the introduction; the following is an overview of these predictors and their levels; the patterns of how genitives are distributed across the levels of previously studied predictors are in line with previous research.

- MODALITY: *spoken* vs. *written*: our sample contains 1,878 genitives from spoken texts, 27.10% of which are *s*-genitives. Of the 2,167 genitives from written texts, only 22.34% are *s*-genitives. Data sparsity permits neither a more fine-grained division of the two modalities into different registers nor an analysis of how MODALITY interacts with other predictors; however, we are not aware of studies of alternation phenomena in which MODALITY interacted with other linguistic predictors in a way that led to a reversal of hypothesised effects anyway.
- ANIMACY (of the possessor): *animate* vs. *collective* vs. *locative* vs. *temporal* vs. *inanimate*. In our sample, the distribution of possessor animacy is as follows: *animate* – 1,028, *collective* – 695, locative – 280, temporal – 184, inanimate – 1,858. The proportion of *s*-genitives is 59.44%, 32.81%, 19.64%, 33.15% and 2.05%, respectively.
- SIBILANCY (of the final phoneme of the possessor): *absent* vs. *present*. In 3,053 cases, there is no final sibilant, but in 992 cases, a final sibilant is present; *s*-genitives are used in 28.15% and 13.53% of the cases, respectively.

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

- DEFINITENESS (of the possessor): *definite* vs. *indefinite*. 2,621 possessors are definite, while 1,424 are not. With definite possessors, the *s*-genitive rate is 31.63%; with indefinite ones, it is 11.52%.
- LENGTHDIFF: the difference of $\log_2$ possessor length minus $\log_2$ possessum length, thus an approximately normally distributed numeric predictor ranging from –4.52 to 4.71 (length is measured in words). In 224 genitives, the possessor and the possessum are equally long; the *s*-genitive is used in 31.70% of these cases. When the possessor is longer (i.e. LENGTHDIFF>0), the *s*-genitive is used in only 14.98% of the cases. If it is shorter (i.e. LENGTHDIFF<0), the *s*-genitive is used in 38.32% of the cases.
- SEMRELATION: *prototypical* vs. *non-prototypical*. In our data, we find 242 prototypical and 3,803 non-prototypical relations. When prototypical, the *s*-genitive is used in 52.90% of the cases and when not, in only 22.75%.

6.2.2 Statistical Evaluation

While this kind of alternation question is one that would prototypically be explored with a generalised linear (mixed-effects) model, we did not proceed along that route. This is due to both the skewed distribution (towards the *of*-genitive) already briefly mentioned in Table 6.1 and the additional fact that the potential random-effects structure looked as if it would become highly problematic: there were a fairly high number of speakers who contributed only few data points to the sample (30% of the data points were by speakers contributing only ten or fewer data points), lowering the chance of proper convergence of our regression models and/or the random effects being particularly relevant. We therefore decided to use an approach based on random forests, an extension of classification and regression trees, here specifically the kind referred to as conditional inference trees (Hothorn et al. 2006) and implemented in R as party::cforest (see also Chapter 7). Random forests add additional layers of randomness to such a tree-based analysis: first, many different conditional inference trees are constructed on different bootstrapped samples of the data. Second, each split in a conditional inference tree is only permitted to choose from a randomly chosen subset of the available predictors rather than all of them. The predictions of the random forest consist of amalgamating the multitude of trees that were generated and their votes for the out-of-bag cases. Typically, the user has to specify only two hyperparameters (i.e. parameters that are defined before a statistical analysis begins

and affect how it is conducted): the number of (randomly chosen) predictors that may be considered at each split of each tree (we left that at the default value of 5) and the number of trees grown (we set that to 2,000).

In order to interpret the results of the random-forest analysis, several strategies are available. One that has been in use especially since Tagliamonte and Baayen (2012) is to i) perform a random-forest analysis on the data, ii) report variable importance scores from the random forest to assess each predictor's importance to the alternation and iii) use a single classification/conditional inference tree on the complete data to visualise the predictors' effects. In this study, we are not following this approach. This is for two main reasons that previous research has ignored. First, the practice of interpreting a random forest – i.e. a set of often 500 or even many more trees on randomly resampled data with different sampled predictors at every split – on the basis of a *single* tree on *all the data* with no resampling is highly problematic and can lead to misinterpretation of the patterns in the data. Second, the way in which random forests are often interpreted – variable importance scores and partial dependency scores – can fail dramatically at representing the nature of the effects in the data faithfully in terms of both over- and underestimated variable importance scores and how predictors interact with each other. Space does not permit a more detailed discussion here; suffice it to say that trees and random forests, which are supposed to be very good at detecting and visualising interactions, are not necessarily as good as they are widely believed to be (see Gries [in press] for more discussion and exemplification and Deshors and Gries [in press] for another English-varieties application).

In order to address all these issues we follow Gries' (in press) recommendations: the first step of our statistical analysis consisted of manually creating a number of new predictors that represent what in a regression model would be interaction predictors, i.e. new variables that embody all combinations of the predictors they consist of:

- all two-way interactions of all predictors with GENDER and VARIETY: GENDER:VARIETY, GENDER:MODALITY, VARIETY:MODALITY, GENDER:ANIMACY, VARIETY:ANIMACY, GENDER:SIBILANCY, VARIETY:SIBILANCY, GENDER:DEFINITENESS, VARIETY:DEFINITENESS, GENDER:LENGTHDIFF, VARIETY:LENGTHDIFF, GENDER:SEMRELATION and VARIETY:SEMRELATION;
- all three-way interactions involving GENDER and VARIETY: VARIETY:GENDER:MODALITY, VARIETY:GENDER:ANIMACY, VARIETY:GENDER:SIBILANCY, VARIETY:GENDER:DEFINITENESS, VARIETY:GENDER:LENGTHDIFF and VARIETY:GENDER:SEMRELATION.

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

These were then added as predictors to a forest of all 2,000 conditional inference trees.

We then evaluated the forest in two ways: first, we computed the forest's overall prediction accuracy, its precision and recall and its *C*-score to determine how well the forest identified structure in our data; second, we computed regular variable importance scores but also an alternative one proposed in Janitza et al. (2013), which is not based on error rates from categorical predictions, which loses important probabilistic information, but in fact on the area under the curve (*AUC*), which does not just rely on categorical predictions but also uses the probabilistic strength of the predictions.

As for evaluating the directions of effects, multiple options are theoretically available, and it does not seem as if there is much of a discussion let alone a consensus yet as to what works best. One could explore effects on the basis of

- the observed percentages of *of-* and *s*-genitives for each level of each predictor (main effects or interaction predictors alike);
- the averages of the predicted percentages of *s*-genitives for every *attested* combination of each level of each predictor;
- the weighted (by frequency of occurrence) averages of the predicted percentages of *s*-genitives for every *theoretically possible* combination of each level of each predictor.

It does not seem that much of the corpus-linguistic literature on random forests topicalises this issue much but, after some consideration, we ultimately decided to go with the last option: while the first approach would be appealing for its simplicity, it has the huge disadvantage that it shows the differences between levels of a predictor but too simplistically, because this would involve levels of a predictor *without* controlling for all other effects or holding all others constant; thus, one would never know to what degree the effect observed for one predictor is also (in part) due to others, which also often leads to exaggerated and anticonservative results. (This, in fact, is the reason why multifactorial regression models should not be summarised with observed means.)

The second approach would be better in that it would be based on predicted, not observed, probabilities and is the logic behind so-called partial dependence statistics/plots (Friedman 2001). However, it seems as if these averages are still suboptimal in how they would not weight predicted probabilities by the frequencies of predictor levels in the data (see Molnar 2018: Section 5.1), thereby – in unbalanced observational

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

data like the present – this might result in upgrading the impact of infrequent combinations and downgrading the impact of frequent ones.

The third approach, while computationally more complex than both the previous ones, seems theoretically most sound and is in fact the logic that underlies Fox's (2003: 1) effect plots for regression models where 'values of other predictors [i.e. all those not currently being computed/visualised] are fixed at typical values: for example, a covariate could be fixed at its mean or median [we do not have any here since, for ease of representation we will factorise LENGTHDIFF], a factor at its proportional distribution in the data'; not only is this much more effective than simple observed results, but this approach also leads to easier-to-interpret results than regression tables and visualises intercepts, main effects and all interactions nicely, which is why we will adopt those plots here. Applied to the genitive alternation, this means that we will – for each combination of predictors of interest (such as VARIETY:GENDER) – inspect the effects of these combinations on genitives with otherwise typical values (in the sense of typical distribution) of remaining covariates, such as ANIMACY, SIBILANCY, etc.

6.3 Results

The random forest of conditional inference trees resulting from the above analysis performed well on the data: the OOB prediction accuracy obtained is 84.5%, which is significantly better than a baseline percentage of the more frequent *of*-genitive (75.5%) and the baseline percentage one would arrive at from random proportional guessing (63%); both $p < 10^{-44}$. Precision and recall for *s*-genitives are not particularly high (71% and 62.1%, respectively), but this is in part due to the class imbalance: precision and recall for the *of*-genitive are much better (88.2% and 91.7%, respectively); with a value of 0.909, the *C*-score for the random forest exceeds the standard threshold value of 0.8.

In terms of variable importance, the top 10 *AUC*-based variable importance values are shown in Table 6.2.

Before we look at some of these predictors' effects, it is instructive to compare this set of variable importance values to those that result from a conditional random forest fitted without interaction predictors, as shown in Table 6.3.

The way in which this is instructive is that the forest without interaction variables does not really encourage the analyst to explore variable combinations/interactions that the forest with interaction variables clearly ranks really highly. For instance and to use the language of regression analysis,

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

Table 6.2 *AUC-based variable importance scores from the conditional inference forest with explicitly coded interaction variables*

Predictor	Var. imp.	Predictor	Var. imp.
ANIMACY	0.0912863	VARIETY:ANIMACY	0.0736966
GENDER:ANIMACY	0.0612236	VARIETY:GENDER:ANIMACY	0.0518154
LENGTHDIFF	0.0375256	GENDER:LENGTHDIFF	0.0134294
VARIETY:GENDER: LENGTHDIFF	0.0118362	VARIETY:LENGTHDIFF	0.0113675
VARIETY:GENDER: DEFINITENESS	0.0061491	VARIETY:GENDER:SIBILANCY	0.0053940

Table 6.3 *AUC-based variable importance scores from the conditional inference forest without explicitly coded interaction variables*

Predictor	Var. imp.	Predictor	Var. imp.
ANIMACY	0.2598940	LENGTHDIFF	0.0797186
SIBILANCY	0.0094320	GENDER	0.0087491
DEFINITENESS	0.0073567	VARIETY	0.0062786
MODALITY	0.0043213	SEMRELATION	0.0038330

while both rankings put ANIMACY first, suggesting to researchers to explore this as a main effect with, for instance, a visual representation of the type in Figure 6.1, the ranking of the forest with interaction variables immediately serves to caution against this, given that ANIMACY appears in interactions with other predictors.

This is relevant here because of how the strongest main effect – ANIMACY – but also the main variables of interest in this analysis – VARIETY and GENDER – are all involved in interactions with a high degree of importance. In fact, the interaction predictors either score more highly than the main effects of which they are made up (see e.g. VARIETY and GENDER's main effects are not among the top ten predictors but they feature in interaction predictors that are) or immediately follow a main effect predictor (see e.g. LENGTHDIFF). Since it is problematic to analyse a random forest with a single tree fitted on all the data (which, if such a tree was not unproblematic, could reveal interactions in the sequence of splits), exploring interaction variables is, therefore, a possible alternative (we will discuss other alternatives below).

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

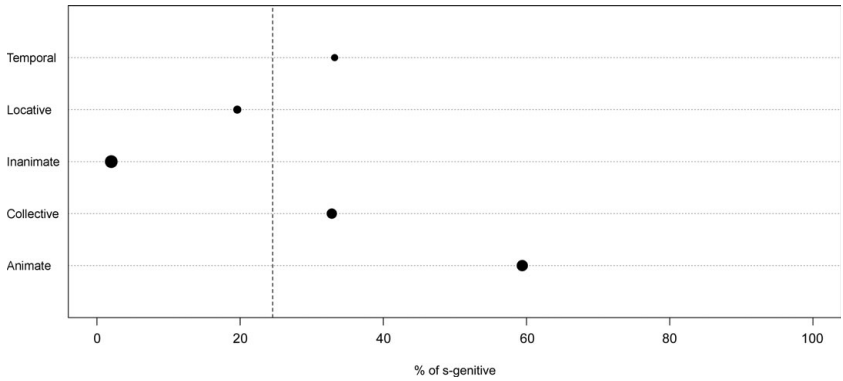

Figure 6.1 Percentages of GENITIVE: *s* for the levels of ANIMACY (vertical dashed line: overall frequency of s-genitives, point sizes are proportional to level frequencies)

In what follows, we will discuss the following effects: VARIETY:GENDER:ANIMACY (Section 6.3.1), VARIETY:GENDER:LENGTHDIFF (Section 6.3.2), VARIETY:GENDER:DEFINITENESS (Section 6.3.3) and VARIETY:GENDER:SIBILANCY (Section 6.3.4); the reason we are focusing on these is that these are the predictors with variable importance scores among the top ten and the ones with the highest order of interactions for every predictor; for instance, ANIMACY has the highest value, but it participates in an interaction in the second most important predictor VARIETY:ANIMACY, but then these two predictors as well as the third most important one, GENDER:ANIMACY, all are involved in the three-way interaction VARIETY:GENDER:ANIMACY, which is therefore the first effect to be discussed.

6.3.1 The Effect of Variety:Gender:Animacy

The first interaction is shown in Figure 6.2. The *y*-axis shows the predicted percentage of *s*-genitives for the combinations of five levels of ANIMACY and two levels of GENDER (abbreviated versions of *male* and *female*, *unknown* is not shown) shown across the *x*-axis. The two varieties are shown as filled circles: BrE in light grey and SriLE in dark grey.

One immediately obvious result is that, with inanimate possessors, there is essentially no difference between varieties and/or genders: *s*-genitives are just very strongly dispreferred (ever so slightly more in BrE). Another fairly strong result is that the interaction is most pronounced for animate

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

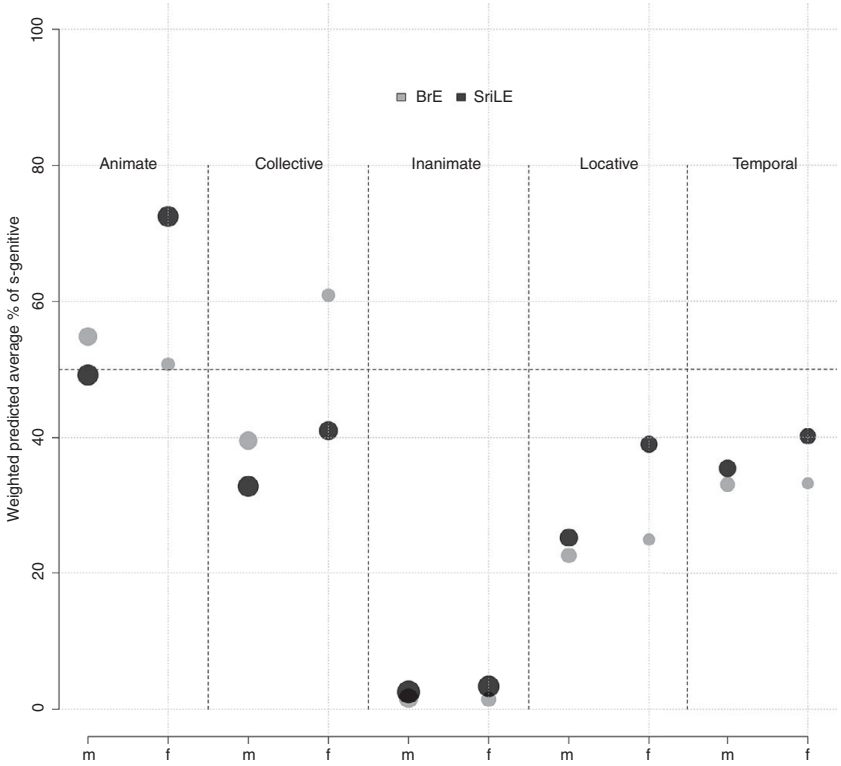

Figure 6.2 The effect of VARIETY:GENDER:ANIMACY on GENITIVE shows that gender differences are most pronounced for genitive choice with animate possessors

possessors: with animate possessors, male speakers of both varieties use the *s*-genitive half the time, but the female SriLE speakers use *s*-genitives much more than the female BrE speakers, who in turn use it less than male speakers; in other words, the female SriLE speakers are exhibiting the canonical/expected patterning more than the female BrE speakers. Finally, we find that collectives behave differently from the other animacy levels: i) *s*-genitives are rarer in SriLE than in BrE and ii) compared to the other combinations, there is a very high percentage of *s*-genitive use among female BrE speakers. While this pattern should not be overinterpreted, given the small number of data points for exactly that combination, female BrE speakers seem more advanced in adopting the expansion of the *s*-genitive to possessors that are lower on the animacy scale (see Wolk et al. 2013 for a diachronic account).

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

Another interesting finding is that, in most combinations of predictors, SriLE speakers use *s*-genitives more than BrE speakers: the dark grey dots are nearly always higher up than the light grey dots. Also, usually, the differences between the varieties are greater with the female speakers (with a slight exception of inanimate possessors, but with these there is nearly a floor effect anyway). With locative and temporal possessors, the results are not that marked: SriLE speakers produce similarly more *s*-genitives than BrE speakers and this is much more noticeable with the female speakers.

6.3.2 The Effect of Variety:Gender:LengthDiff

The second relevant interaction is represented in Figure 6.3.

The most obvious result here is the very strong expected main effect of LENGTHDIFF: as the possessor becomes longer relative to the possessum, the *s*-genitive becomes less and less preferred across all combinations of VARIETY and GENDER. The two outer quarters with the rarer situations of big discrepancies between possessor and possessum lengths show less in terms of differences between genders and varieties, but the middle two quarters of the plot, where most of the cases are located, are more interesting: they show that the differences are small between varieties but more noteworthy between men and women because women simply use many more *s*-genitives than men in both varieties.

6.3.3 The Effect of Variety:Gender:Definiteness

The third relevant interaction is represented in Figure 6.4. We can see a main effect such that definite possessors (left panel) lead to higher numbers of *s*-genitives than indefinite possessors (right panel), as might be expected. Also, there is another main effect such that, with one exception, women use more *s*-genitives than men. However, these main effects are qualified by crossover effects. First, the one exception of the main effect just mentioned: female BrE speakers use *s*-genitives less than male BrE speakers, but only with indefinite possessors. Second, the differences between men and women are more pronounced in SriLE than in BrE and that is especially true for definite possessors, where female SriLE speakers exhibit a much higher proportion of *s*-genitives than any other combination, which indicates that Sri Lankan women seem to react more to the grammatical cue of DEFINITENESS than BrE women: for them, the

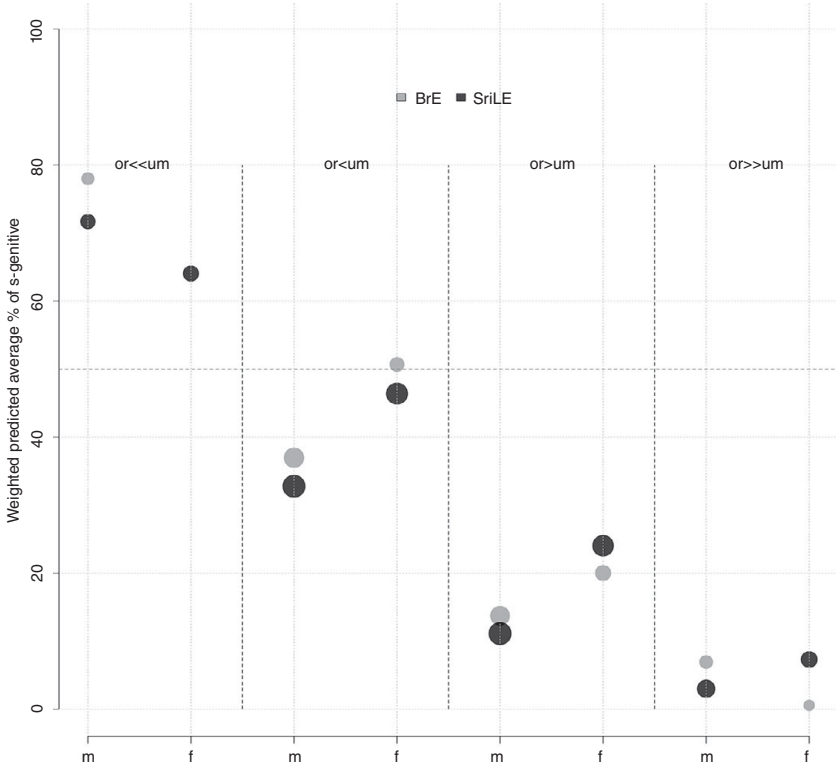

Figure 6.3 The effect of VARIETY:GENDER:LENGTHDIFF on GENITIVE (LENGTHDIFF is factorised into four levels for expository purposes) is subject to stronger gender differences in SriLE (dark grey) than in BrE (light grey)

difference from definite to indefinite is bigger – male speakers show less of an impact there.

6.3.4 *The Effect of Variety:Gender:Sibilancy*

The final interaction to be discussed is shown in Figure 6.5. There is the overall expected main effect of SIBILANCY, according to which SIBILANCY: *present* (right panel) should reduce the presence of the *s*-genitive and if there is a sibilant, then both genders use *s*-genitives correspondingly rarely. However, female speakers react more to SIBILANCY than men (and especially so in SriLE) and SriLE speakers react more to SIBILANCY than BrE speakers (and especially so the female speakers).

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

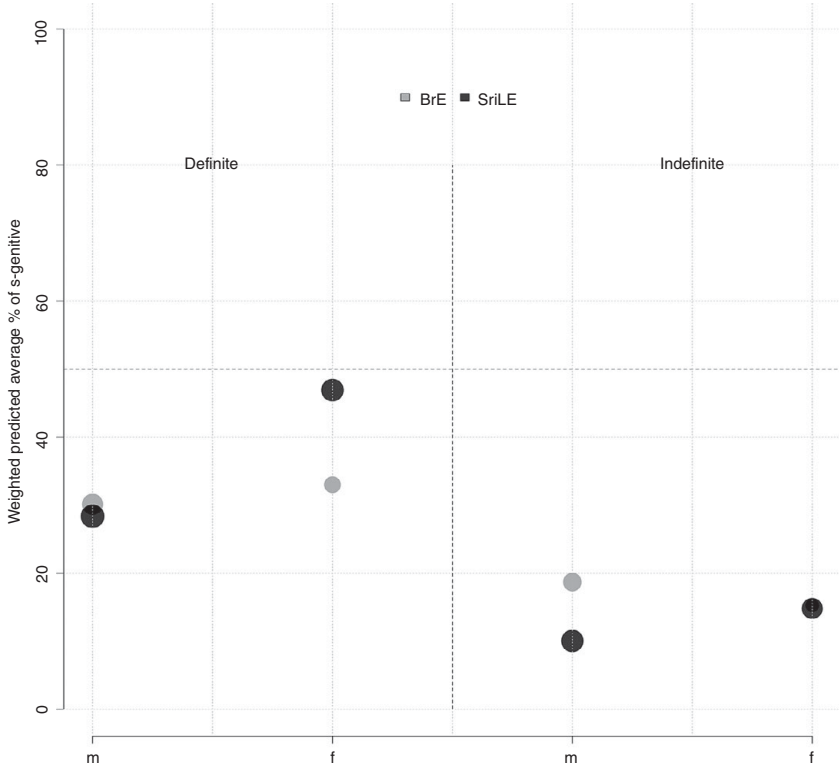

Figure 6.4 The effect of VARIETY:GENDER:DEFINITENESS on GENITIVE shows that SriLE speakers and females are more sensitive to the definiteness constraint

6.4 Discussion and Concluding Remarks

6.4.1 Implications for the Genitive Alternation

On the whole and on a general linguistic level, our results are largely compatible with previous studies: we find a strong effect of (possessor) ANIMACY that is in line with previous findings; the same is true of the strong effect of LENGTHDIFF. While weaker, the effects of DEFINITENESS and SIBILANCY do not contradict prior research either. Reassuringly, we find that a factor such as LENGTHDIFF seems to be interacting less with GENDER, which makes sense since one would expect male and female speakers to have a very similar cognitive architecture and, thus, react similarly to the higher degree of processing pressure that arises from high

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

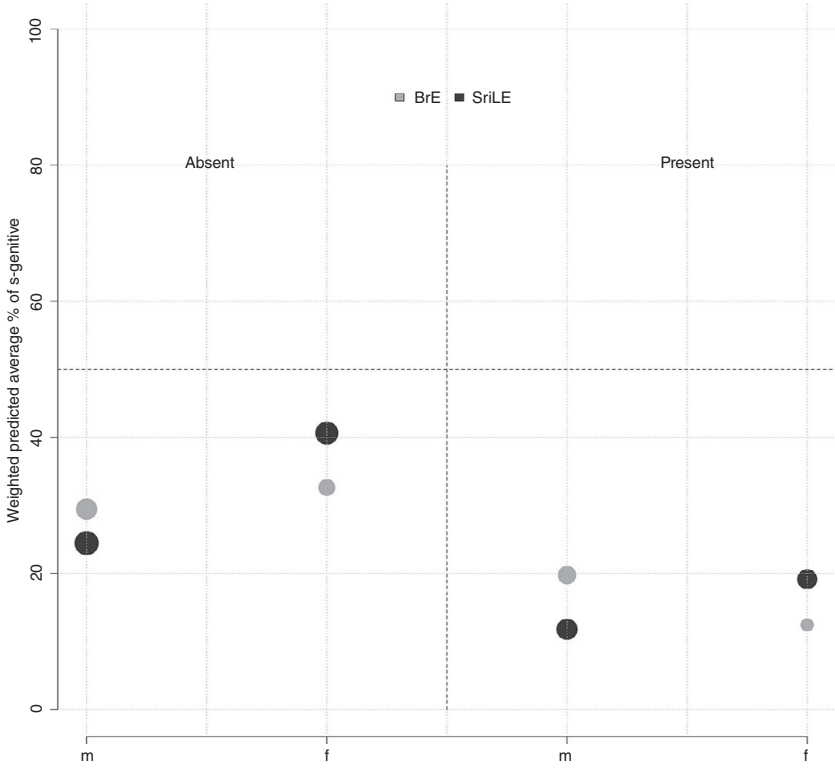

Figure 6.5 The effect of VARIETY:GENDER:SIBILANCY on GENITIVE shows particular sensitivity of SriLE speakers and females

length differences. However, there are also more noteworthy differences between the genders.

The results displayed in Figures 6.2–6.5 show that the language-external factors GENDER and VARIETY moderate the effects of language-internal constraints on genitive choice. As stated in the introduction, this conforms to our expectations derived from recent research, which uncovered how VARIETY modulates the effect of ANIMACY (e.g. Heller et al. 2017b). However, our results also go against and beyond the expected in that i) the effect strength of the ANIMACY constraint across the varieties differs from expectations and ii) the effect of GENDER on genitive choice had not yet been investigated in a multifactorial design.

Regarding i), we find that the effect of ANIMACY – the most important language-internal predictor of genitive choice – more precisely, the

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

difference between animate and inanimate possessors, is stronger in SriLE than in BrE. This is unexpected because recent research (Heller et al. 2017b; Heller 2018) found the effect of ANIMACY to be weaker in Outer-Circle varieties (i.e. Hong Kong, Indian, Jamaican, Philippine and Singapore English) than in Inner-Circle varieties (i.e. British, Canadian, Irish and New Zealand English). Since SriLE arguably qualifies as an Outer-Circle variety (e.g. Schneider 2011), it was expected that the effect of ANIMACY would be weaker in SriLE. Counter-intuitively, however, Figure 6.2 shows bigger differences between the animate condition and the inanimate condition in SriLE.

Regarding ii), we present first findings on how GENDER enters the equation: GENDER moderates the effects of length difference and variety and also further qualifies the interactions between VARIETY and ANIMACY (see the previous paragraph) as well as between VARIETY and DEFINITENESS. First, there appears to be a slight gender difference in the effect of length in that males show up as more sensitive to the condition in which the possessum is much longer than the possessor (see the leftmost panel in Figure 6.3). In this condition, males use the *s*-genitive more frequently. In all other conditions, females seem to be more drawn to the *s*-genitive (see other panels in Figure 6.3). Closer inspection of these differences reveals that, on aggregate, males respond to length difference in a more categorical fashion: while females respond to length difference fairly linearly (i.e. the stronger the cue, the stronger their reaction), males mostly default to using the *of*-genitive, but as soon as possessum length outgrows possessor length by a certain threshold, they prefer the *s*-genitive and, within this extreme range, do so even more than females. GENDER also interacts with VARIETY since most of the time we see that light grey m/f dots are closer to each other than the dark grey m/f dots (Figures 6.2–6.5). Depending on which predictor we interpret as focal, we can describe the observed preference for the *s*-genitive in two ways: either 1) in SriLE, we find a stronger effect of this preference in females, which presumes a stronger gender difference in SriLE, or 2) in females, we find a stronger effect of this preference in SriLE, which presumes that cross-varietal differences emerge more strongly in female language use. In other words, Sri Lankan females use the *s*-genitive more than Sri Lankan males and gender differences are more pronounced in SriLE. Finally, GENDER mediates the VARIETY-ANIMACY and VARIETY-DEFINITENESS interactions: females appear to be more sensitive to both. Once GENDER is taken into account, we see that the stronger effect of ANIMACY in SriLE goes back to female language users only; Sri Lankan males, on the other hand,

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

behave fairly similarly to BrE speakers. A similar effect emerges in the possessor DEFINITENESS constraint. SriLE-speaking females respond to DEFINITENESS more strongly than SriLE-speaking males or speakers of BrE. Thus, Sri Lankan females show higher sensitivity to both ANIMACY and DEFINITENESS constraints.

How can we make sense of these patterns? Although there are obvious limitations to our study (e.g. partly exploratory design and limited sample size), it seems reasonable to propose a contact-linguistic explanation of our findings. Since our study covers uncharted territory by focusing on the role of GENDER in the probabilistic grammar of English genitive choice, we cannot inform any World Englishes models (and vice versa) because these models do not make predictions along the lines of GENDER (e.g. models by Kachru, McArthur, Schneider, or more recently by Mair or Buschfeld and Kautzsch). Therefore, and based on previous studies such as Brunner (2014), we turn to a more specific contact-linguistic explanation of the gender difference in the strength of the possessor ANIMACY constraint.

The stronger inclinations of (female) SriLE speakers to use the *s*-genitive might be caused by a transfer of structures in Sinhala, Sri Lanka's most prevalent native language. This transfer might work in two ways – directly and/or in a more abstract way. In Sinhala, the possessor always precedes the possessum (see Chandralal 2010: 10), which corresponds to the *s*-genitive. This constituent ordering might carry over directly from Sinhala to English, equivalent to the transfer found by Brunner (2014), who observed that noun phrase modification patterns in Singapore English and Kenyan English correlate with preferences in the countries' respective native languages. There could also be a more abstract transfer of cue strength. Rosenbach (2017) showed that this is indeed the case with genitive choice in the L2 English of Afrikaans speakers. However, the relation between Sinhala and SriLE is different because there is no genitive alternation in Sinhala – the transfer of the ANIMACY constraint could thus only be plausible on a more abstract level. We propose that it might be the high salience of the constraint in Sinhala that carries over to English. Sinhala is special in that it has a different inflectional morphology for animate/inanimate and definite/indefinite nouns (see Chandralal 2010: 45). The distinctions are thus ubiquitous. Because of the high salience of the ANIMACY and DEFINITENESS constraints, speakers could more easily pick up on the constraints in English and also use these constraints in a more categorical fashion, resulting in higher usage frequencies of the *s*-genitive with animate/definite possessors. High salience might also cause an overcorrective use of the animacy rule (i.e. use the *s*-genitive with animate

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

possessors); this tendency to overcorrect might be further enhanced by the perceived high status of English in postcolonial societies (Schneider 2007).

However, it remains unclear why these cross-varietal differences in the effects of ANIMACY and DEFINITENESS mostly rely on preferences found in Sri Lankan females. A tentative search for explanations might include societal factors, such as gender equality. According to the United Nations Development Programme (2014), the societies of Sri Lanka and Great Britain are vastly different in terms of labour market participation. In Sri Lanka, only 35% of women participate in the labour market (males: 76.4%), whereas in the United Kingdom, 55.7% of women participate (males: 68.8%) (see United Nations Development Programme 2014: 172–3). Higher participation in the workforce might require more use of English and wider social networks, which might explain the more BrE-like patterns of Sri Lankan males. Sri Lankan women's relative absence in the labour market might further explain why we observe less convergence between the genders than we do in Britain.[2] Lower participation in the workforce, however, is likely to be associated with less use of English and narrower social circles, which might pave the way for more influence from Sinhala in the English of Sri Lankan females.

6.4.2 *Methodological Implications*

This study has methodological implications as well. First, it is one of the first studies following Gries' (in press) recommendations regarding classification trees and random forests (see also Deshors and Gries in press). To avoid using summarising a random forest on the basis of a single tree and an interpretation biased towards main effects, we used Gries' new random-forest protocol, which involves including interaction predictors and *AUC*-based variable importance scores to determine whether interactions are relevant, too, and which predictors (main effects or interactions) to discuss. We believe that the comparison between Figures 6.1 and 6.2 clearly shows that main effects, even if they are the most highly ranked predictors of the analysis, underestimate, or at least do not at all highlight, the complexity of a dataset (and, again, a single tree cannot be used reliably to determine the interaction structure in a whole random forest).

Second, we are the first to implement an alternative to partial dependency scores, which are not yet readily implemented for the experimental

[2] We would like to thank Reviewer 2 for this suggestion.

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

forests of conditional inference trees coming with the PARTY package for R, namely, an exploration of predicted probabilities that mirror those of the widely used effects plots à la Fox (2003) for regression modelling. We think that this is a useful way to proceed given how much using observed percentages can lead to anticonservative overestimates of effects especially in infrequent (combinations of) levels of predictors and in cases with moderate to high collinearity in the data, in which case the effect that a plot returns for one predictor will also contain the effect of many other correlated predictor values.

To see that effect in the present dataset, consider Figure 6.6, which is a somewhat complex extension of Figure 6.3. Specifically, the filled circles are the points from Figure 6.3 (with one default size), but,

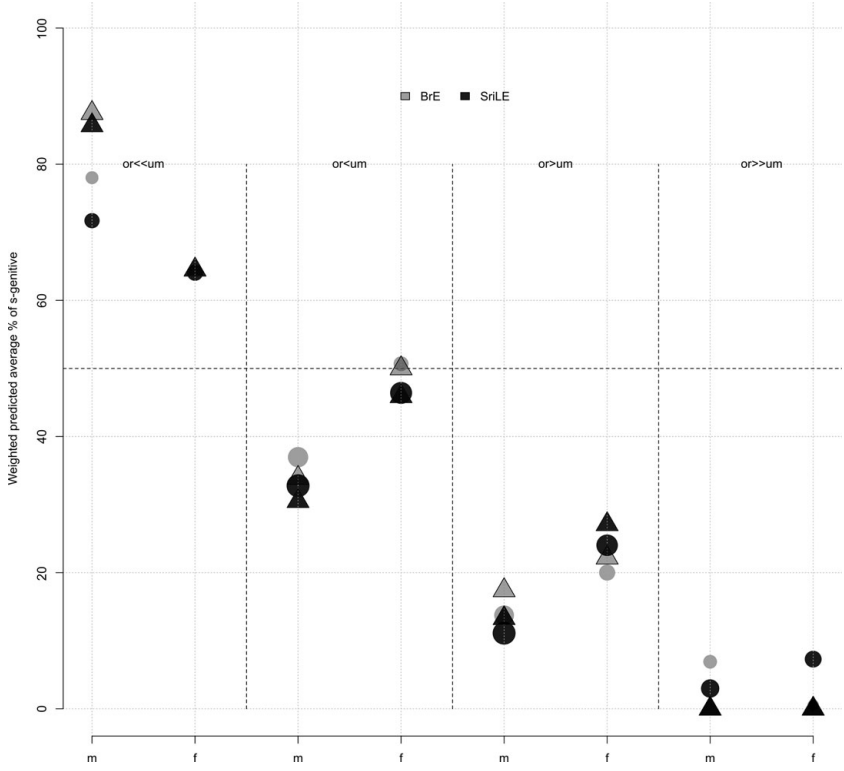

Figure 6.6 The effect of VARIETY:GENDER:LENGTHDIFF on GENITIVE (circles) vs. observed percentages of *s*-genitives for VARIETY:GENDER:LENGTHDIFF (triangles)

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

crucially, the light grey and dark grey triangles represent the simple observed percentages of *s*-genitives in BrE and SriLE; in other words, the further light grey and dark grey triangles are away from light grey and dark grey circles, respectively, the more the simple observed percentages differ from the effects-type predictions computed here that control for other predictors.

The interpretation is relatively straightforward: in the middle two quarters of the plot, which represent length differences that are fairly frequent in the data, the two estimation methods lead to similar results: the percentages are a bit off and sometimes the ratio of BrE to SriLE is incorrect in the observed percentages display, but the differences are not too dramatic. However, in the outer two quarters of the plot, which represent much fewer cases – because huge discrepancies between possessor and possessum length are rarer – the triangles/observed percentages exaggerate the effects much more.

For instance, in the rightmost quarter, the observed percentages are all 0 whereas the effects predictions are (sometimes quite a bit) higher. It is important to realise that while there are no *s*-genitives when the possessor is much longer than the possessum, which the triangles represent, the triangles/observed percentages are still not a good guide towards understanding the effect of LENGTHDIFF. This is because, for instance, these cases where the possessor is much longer than the possessum also have

- a much higher number of indefinite possessors, which also favour *of*-genitives;
- a much higher number of possessors with final sibilants, which also favour *of*-genitives.

Thus, the triangle positions at $y = 0$ reflect *multiple* variables' effects, not just, like the figure caption would have one believe, (VARIETY:GENDER:) LENGTHDIFF.

The same is true of the or<<um cases on the left, which involve many more definite and non-sibilant-ending possessors than overall so here, too, the triangles/observed percentages overestimate what the graph has a reader attribute to (VARIETY:GENDER:)LENGTHDIFF. Also, the effects-like computation for random-forest predictions we pioneer here is, we believe, a nice way of extending a tried-and-true logic from regression modelling to random forests to lead to a better understanding of their effects (one that is comparable to the use of global surrogate models for random forests, see Gries in press).

https://doi.org/10.1017/9781108696739.006 Published online by Cambridge University Press

6.4.3 Where to Go from Here

While our study is an exercise in World Englishes scholarship, its results do not straightforwardly inform popular models in the field because World Englishes models do not make predictions on our subject matter. This study is concerned with an abstract syntactic alternation and the slight probabilistic differences in its conditioning factors and in particular their interactions with GENDER. Current models of World Englishes do not make predictions about these things; in fact, before our study, there has hardly been any indication that male and female speakers make different genitive choices. In this sense, our findings comply equally with all World Englishes models.

That being said, this study has shown that GENDER is an important determinant of probabilistic grammatical choice-making and we do offer a potential explanation for our strongest finding – the gender difference in the strength of the possessor animacy constraint – by referring to a possible L1 transfer. While L1 transfer, again, is compatible with, and thus does not distinguish between, basically all popular models of World Englishes (e.g. by Kachru, McArthur, Schneider, or more recently by Mair or Buschfeld and Kautzsch), seeing these gender differences is instructive both in and of itself but also in terms of how GENDER qualifies other interactions and in terms of how it may force the analyst to face the sociocultural realities on the ground and how they impact linguistic choices and language change. Further studies of genitive choice should thus do their best to take this influence into account.

In order to understand the role of gender on grammatical alternations more fully, researchers should also look at its effects in a larger scope. This might involve 1) looking at genitive choice in more than two varieties or 2) looking at the effect of gender in other positional alternations such as the dative alternation or the particle placement alternation. To facilitate this, we recommend the use of additional ICE metadata provided by Martin Schweinberger (see above) or Beke Hansen (Hansen 2017).

Further, we see potential of further research in the field of contact-induced probabilistic differences in grammatical alternations across varieties. Although previous studies have suggested a transfer of constituent order (Brunner 2014) and probabilistic weights (Rosenbach 2017), contrasting findings that show opposite patterns exist as well (Heller 2018). Further, the present study suggests that high salience of certain distinctions might already spark probabilistic differences in English, a hypothesis that remains to be tested. This might be achieved by inspecting genitive choice

across a range of varieties with differing degrees of salience and different probabilistic weights of predictors like ANIMACY and DEFINITENESS in the respective countries' L1s.

Finally, we do feel that the statistical methodology we applied here merits more attention and future use. While mixed-effects modelling in particular has been taking much of linguistics by storm, its applicability to observational data is still often quite difficult so it is understandable that alternatives such as tree-based methods and/or random forests are becoming prominent alternatives. However, as alluded to above, if only briefly, there are scenarios in which the deceptive simplicity of these methods is counterproductive and hides some of the interesting variability in the data – exploring interactions and visualising effects are extremely rare in random-forest studies and we hope to have shown how and why this matters and what the discipline has to gain from such steps; the main effects of many phenomena are already well understood so the ability to add interactions with, for instance, speaker-level effects or other language-external factors is one of the things we need to move things to the next level.

REFERENCES

Altenberg, Bengt. 1982. *The Genitive v. the of-Construction: A Study of Syntactic Variation in 17th Century English*. Lund: Gleerup.

Behaghel, Otto. 1909. 'Beziehungen zwischen Umfang und Reihenfolge von Satzgliedern', *Indogermanische Forschungen* 25: 110–42.

Bernaisch, Tobias. 2012. 'Attitudes towards Englishes in Sri Lanka', *World Englishes* 31(3): 279–91.

2015. *The Lexis and Lexicogrammar of Sri Lankan English*. Amsterdam: John Benjamins.

Bock, J. Kathryn. 1982. 'Toward a cognitive psychology of syntax: information processing contributions to sentence formulation', *Psychological Review* 89 (1): 1–47.

Brunner, Thomas. 2014. 'Structural nativization, typology and complexity: noun phrase structures in British, Kenyan and Singaporean English', *English Language and Linguistics* 18(1): 23–48.

Chandralal, Dileep. 2010. *Sinhala*. Amsterdam: John Benjamins.

Deshors, Sandra C. and Gries, Stefan Th. in press. 'Mandative subjunctive vs. should in world Englishes: a new take on an old alternation', *Corpora*.

Fox, John. 2003. 'Effect displays in R for generalised linear models', *Journal of Statistical Software* 8(15): 1–27.

Friedman, Jerome H. 2001. 'Greedy function approximation: a gradient boosting machine', *Annals of Statistics* 29(5): 1189–232.

Grafmiller, Jason. 2014. 'Variation in English genitives across modality and genres', *English Language and Linguistics* 18(3): 471–96.
Greenbaum, Sidney. 1996. 'Introducing ICE'. In Sidney Greenbaum, ed. *Comparing English Worldwide: The International Corpus of English*. Oxford: Clarendon, 3–12.
Gries, Stefan Th. in press. 'On classification trees and random forests in corpus linguistics: some words of caution and suggestions for improvement', *Corpus Linguistics and Linguistic Theory*.
Gunesekera, Manique. 2005. *The Postcolonial Identity of Sri Lankan English*. Colombo: Vijitha Yapa Publications.
Hansen, Beke. 2017. 'The ICE metadata and the study of Hong Kong English', *World Englishes* 36(3): 471–86.
Heller, Benedikt. 2018. *Stability and Fluidity in Syntactic Variation World-Wide: The Genitive Alternation across Varieties of English* (Unpublished doctoral dissertation). KU Leuven, Leuven, Belgium.
Heller, Benedikt, Bernaisch, Tobias and Gries, Stefan Th. 2017a. 'Empirical perspectives on two potential epicenters: the genitive alternation in Asian Englishes', *ICAME Journal* 41: 111–44.
Heller, Benedikt, Szmrecsanyi, Benedikt and Grafmiller, Jason. 2017b. 'Stability and fluidity in syntactic variation world-wide', *Journal of English Linguistics* 45(1): 3–27.
Hothorn, Torsten, Hornik, Kurt and Zeileis, Achim. 2006. 'Unbiased recursive partitioning: a conditional inference framework', *Journal of Computational and Graphical Statistics* 15(3): 651–74.
Janitza, Silke, Strobl, Carolin and Boulesteix, Anne-Laure. 2013. 'An AUC-based permutation variable importance measure for random forests', *BMC Bioinformatics* 14: 119.
Kachru, Braj B. 1985. 'Standards, codification and sociolinguistic realism: the English language in the outer circle'. In Randolph Quirk and Henry G. Widdowson, eds. *English in the World: Teaching and Learning the Language and Literatures*. Cambridge University Press, 11–30.
Kendall, Tyler, Bresnan, Joan and van Herk, Gerard. 2011. 'The dative alternation in African American English: researching syntactic variation and change across sociolinguistic datasets', *Corpus Linguistics and Linguistic Theory* 7(2): 229–44.
Künstler, Viktoria, Mendis, Dushyanthi and Mukherjee, Joybrato. 2009. 'English in Sri Lanka: language functions and speaker attitudes', *International Journal of English Studies* 20(2): 57–74.
Labov, William. 1972. *Sociolinguistic Patterns*. Philadelphia: University of Philadelphia Press.
MacDonald, Maryellen C. 2013. 'How language production shapes language form and comprehension', *Frontiers in Psychology* 4: 1–16.
Meyler, Michael. 2007. *A Dictionary of Sri Lankan English*. Colombo: Mirisgala.
Molnar, Christoph. 2018. 'Interpretable machine learning: a guide for making black box models explainable' (Manuscript). Retrieved from https://christophm.github.io/interpretable-ml-book/index.html.

Mukherjee, Joybrato. 2008. 'Sri Lankan English: evolutionary status and epicentral influence from Indian English'. In Klaus Stierstorfer, ed. *Proceedings / Anglistentag 2007 Münster*. Trier: Wissenschaftlicher Verlag Trier, 359–68.

Revis, Melanie and Bernaisch, Tobias. 2020. 'Pragmatic nativisation in Asian Englishes? evidence from filled and unfilled pauses', *World Englishes* 39(1): 135–153.

Rosenbach, Anette. 2002. *Genitive Variation in English: Conceptual Factors in Synchronic and Diachronic Studies*. Berlin: Mouton de Gruyter.

2005. 'Animacy versus weight as determinants of grammatical variation in English', *Language* 81(3): 613–44.

2014. 'English genitive variation – the state of the art', *English Language and Linguistics* 18(2): 215–62.

2017. 'Constraints in contact: animacy in English and Afrikaans genitive variation – a cross-linguistic perspective', *Glossa: A Journal of General Linguistics* 2(1): 72. 1–21.

Schneider, Edgar W. 2007. *Postcolonial English: Varieties around the World*. Cambridge University Press.

2011. *English around the World: An Introduction*. Cambridge University Press.

Szmrecsanyi, Benedikt. 2017. 'Variationist sociolinguistics and corpus-based variationist linguistics: overlap and crosspollination potential', *Canadian Journal of Linguistics/Revue canadienne de linguistique* 62(4): 1–17.

Szmrecsanyi, Benedikt and Hinrichs, Lars. 2008. 'Probabilistic determinants of genitive variation in spoken and written English: a multivariate comparison across time, space, and genres'. In Terttu Nevalainen, Irma Taavitsainen, Päivi Pahta and Minna Korhonen, eds. *The Dynamics of Linguistic Variation: Corpus Evidence on English Past and Present*. Amsterdam: John Benjamins, 291–309.

Szmrecsanyi, Benedikt, Biber, Douglas, Egbert, Jesse, and Franco, Karlien. 2016. 'Toward more accountability: modeling ternary genitive variation in Late Modern English', *Language Variation and Change* 28(1): 1–29.

Szmrecsanyi, Benedikt, Grafmiller, Jason, Bresnan, Joan, Rosenbach, Anette, Tagliamonte, Sali A. and Todd, Simon. 2017. 'Spoken syntax in a comparative perspective: the dative and genitive alternation in varieties of English', *Glossa: a Journal of General Linguistics* 2(1): 86. 1–27.

Tagliamonte, Sali A. and 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): 135–78.

United Nations Development Programme. 2014. 'Human development report 2014: sustaining human progress: reducing vulnerabilities and building resilience' (Manuscript). Retrieved from http://hdr.undp.org/sites/default/files/hdr14-report-en-1.pdf.

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