

1
2
3
4 **Agreement syncretization and the loss of null subjects:**
5 **quantificational models for Medieval French**
6

7 ALEXANDRA SIMONENKO

8 *FWO/UGhent*

9
10 BENOIT CRABBÉ

11 *LLF, CNRS/Paris Diderot, USPC*

12
13 SOPHIE PRÉVOST

14 *LATTICE, CNRS/ENS/Paris Sorbonne Nouvelle, USPC & PSL*
15
16
17

18 ABSTRACT

19 This paper examines the nature of the dependency between the availability of null
20 subjects and the “richness” of verbal subject agreement, known as Taraldsen’s
21 Generalisation (Adams, 1987; Rizzi, 1986; Roberts, 2014; Taraldsen, 1980). We
22 present a corpus-based quantitative model of the syncretization of verbal subject
23 agreement spanning the Medieval French period and evaluate two hypotheses
24 relating agreement and null subjects: one relating the two as reflexes of the same
25 grammatical property and a variational learning-based hypothesis whereby
26 phonology-driven syncretization of agreement marking creates a learning bias
27 against the null subject grammar. We show that only the latter approach has the
28 potential to reconcile the intuition behind Taraldsen’s Generalisation with the fact
29 that it has proven nontrivial to formulate the notion of agreement richness in a way
30 that would unequivocally predict whether a language has null subjects.

31 This paper examines the nature of the relation between the availability of null
32 subjects and the “richness” of verbal subject agreement, known as Taraldsen’s
33 Generalisation (Adams, 1987; Rizzi, 1986; Taraldsen, 1980), from the point of
34 view of grammar change in Medieval French. The original generalization based
35 on synchronic observations states that a language having sufficiently discriminating,
36 or nonsyncretic, subject agreement entails the possibility of nonexpression of
37
38

39 This work has been supported by the Labex EFL and Research Foundation Flanders. The authors are
40 very thankful to three anonymous reviewers of *Language Variation and Change* for extremely
41 helpful comments. We would like to express our deep gratitude to Yves Charles Morin, Henri
42 Kauhanen, George Walkden, Paul Hirschbühler, Philippe Schlenker, and Hedde Zeijlstra for
43 discussions and suggestions. The project has benefited from the feedback from the audiences at
44 WCCFL 34, DiGS 2016, XLanS: Triggers of language change in the Language Sciences, workshop
45 Linguistic Knowledge & Patterns of Variation, Texts, Tools, and workshop Methods in Digital
Classics and Medieval Studies, seminars at Institut Jean Nicod, Université Diderot Paris 7, and the
University of Manchester.

46 subjects. In terms of diachronic developments, it was argued that there is a causal
 47 relation between the loss of nonsyncretic subject agreement and the emergence of
 48 obligatory subject pronouns (e.g., Ewert, 1943; Vennemann, 1975:298), the
 49 underlying intuition being that overt subjects take over the role of identifying the
 50 subject’s person which can no longer be fulfilled by verbal inflection due to its
 51 phonological erosion. Haspelmath (1999:14) said that “... in languages that are
 52 losing their rich subject agreement morphology on the verb ... speakers will
 53 increasingly tend to choose the option of using the personal pronoun, because
 54 the verbal agreement does not provide the information required for referent
 55 identification in a sufficiently robust way.”

56 This diachronic scenario, however, was questioned for Medieval French on the
 57 grounds of an apparent temporal lag between the loss of null subjects and loss of
 58 agreement (e.g., Roberts, 2014; Schøsler, 2002). However, opposite assumptions
 59 have been made about the temporal sequence of the two changes, due to the
 60 unavailability of a systematic quantitative study of syncretization. We present a
 61 corpus-based study spanning the Medieval French period to evaluate two
 62 hypotheses. First, we test the predictions generated by the hypothesis that null
 63 subjects and nonsyncretic agreement exponents are related at the clause level,
 64 both being dependent on the same functional head. The second hypothesis we
 65 explore is based on Yang’s (2002) variational learning model whereby the
 66 agreement exponents and subject expression are not strictly connected at
 67 the clause level. Instead, in the process of language learning (possibly over the
 68 speaker’s lifespan) syncretic endings create a bias against the null subject
 69 grammar, which eventually drives it to extinction.

72 NULL SUBJECTS AND SUBJECT AGREEMENT IN FRENCH

73
 74 Our estimates are based on the corpus of the project “Modéliser le changement: les
 75 voies du français” (MCVF) and Penn Supplement to MCVF (2010), which together
 76 include 35 syntactically parsed texts ($n \approx 1$ million words [Appendix B]). On the
 77 assumption that null subjects correspond to phonologically null personal
 78 pronominal elements, observations about the emergence of overt subjects are
 79 given here as the estimated probability of overt personal pronominal subjects
 80 against null subjects, with demonstrative, nominal, and other kinds of overt
 81 subjects being excluded from consideration. The assumption is warranted by the
 82 fact that the rate of overt subjects that are not personal pronouns stays the same
 83 throughout the Medieval period, whereas the rate of overt pronominal subject
 84 increases and the rate of null subjects decreases in a dramatic fashion.
 85 Furthermore, null subjects in pro-drop languages and overt pronominal subjects
 86 in obligatory subject languages are said to be distributionally equivalent (e.g.,
 87 Hirschbühler, 1992).

88 Our dataset includes all finite clauses with either an overt pronominal or null
 89 subject ($n = 56615$), excluding imperatives, subject relatives, and wh-questions
 90 targeting subjects because of their idiosyncratic subject syntax.¹ We also

91 excluded all coordinated clauses introduced by the coordinating conjunction *et* and
 92 the conjunctive adverb *si*, since those license subject ellipsis throughout the
 93 Medieval period. Although these connectives are sometimes used even when
 94 there is no potential antecedent in the preceding clause, we take the nearly stable
 95 rate of subject omission with *et* and *si* (see Appendix C, Figure 1C) to mean that
 96 there are few true subject omission environments with these connectives. Subject
 97 ellipsis under coordination with *et* is still allowed in Modern French, while *si*
 98 itself fell out of use as a conjunctive adverb.² The nonexpression of referential
 99 subjects occurred in Medieval French, and Old French in particular, in contexts
 100 where their expression would be obligatory in Modern French (e.g., Foulet
 101 [1928] and much literature since). During the Medieval period nonexpression
 102 became more and more rare, for both main and subordinate clauses, as seen in
 103 Figure 1. As has been noted before, subordinate clauses favor overt subjects
 104 more than main clauses (e.g., Foulet, 1928; Franzen, 1939; Hirschbühler, 1992;
 105 Roberts, 2014; Vance, 1997; Zimmermann, 2014; among others), though null
 106 subjects can be found in all types of subordinates (Fontaine, 1985; Hirschbühler
 107 & Junker, 1988; Kaiser, 2009; Prévost, 2018; Roberts, 1993).

108 *Subject agreement syncretization*

109 French went from a language characterized by nonsyncretic agreement inherited
 110 from Late Latin to a language with a largely syncretic agreement paradigm
 111 (Bettens, 2015; Buridant, 2000; De Jong, 2006; Dees, Meilink, van Reenen-
 112 Stein, & van Reenen, 1980; Foulet, 1935; Marchello-Nizia, 1992; Morin, 2001).
 113 We can say that there is no systematic person marking on the verb in Modern
 114 French, and the only subject agreement feature present is number.³ In contrast,
 115 as evidenced by the system of rhymes used in Old French versification, verbal
 116 paradigms had a much less syncretic nature during that period (e.g., Bettens, 2015).
 117

118 Overall, there are three classes of changes that resulted in syncretism, namely,
 119 the drop of the final *-t* after vowels, *e*-insertion, and *s*-insertion. The first two
 120 changes can be seen as related on the hypothesis of Dees et al. (1980) and van
 121 Reenen and Schøsler (1987) that *e*-insertion was a compensatory process
 122 “keeping” root consonants from the final position where they would have fallen.
 123 As we will see below, they are also much closer in time and in how they spread
 124 to each other than to the third one, *s*-insertion. Appendix A details the main
 125 changes in verbal agreement, by verb Group and tense-aspect form. These are:
 126

127 A. innovative final *-e*: 1st person, Group I, present indicative & subjunctive

128 The use of the ending *-e* instead of zero for the 1st person singular subjects with
 129 Group I verbs began in the 12th century, and, by the beginning of the 15th
 130 century, generalized onto the roots ending in a consonant, the zero ending
 131 lingering for longer with stems ending in a vowel (Marchello-Nizia, 1992:200).
 132 A handful of verbs whose stems etymologically ended in *-e*, such as *monstre-r*
 133 ‘to show’ were not affected by this change.

134 B. innovative final *-e* (i.e., becoming final as a result of the drop of *-t*): 3rd person, 135 Group I, present indicative & subjunctive




FIGURE 1. Overt pronominal subjects in main and subordinate clauses ($n = 76150$).

Note: Absolute numbers of null and overt pronominal subjects for each text are given in Appendix B. We use frequencies to estimate probabilities.

The emergence of final *-e* as a consequence of the disappearance of the final *-t* in the context of the 3rd person singular subjects is generally considered to predate the changes in the 1st person singular contexts.

C. innovative final *-e*: 3rd person, Group II, present subjunctive

The alternation between *-et* and an innovative *-e* as the endings of the 3rd person singular present subjunctive in Group II also resulted in syncretism.

D. innovative final *-a*: 3rd person, Group I, preterite & future indicative (did not result in syncretism)

E. innovative final *-a*: 3rd person, Group II future indicative (did not result in syncretism)

F. innovative final \emptyset : 3rd person, Group II, preterite⁶

In Group II, the ending *-t* alternated with zero in the context of the 3rd person singular in preterite. This case is special in that the innovative zero ending was on the rise up until the mid-14th century when it suddenly went into a sharp decline, the old ending reinstalling itself completely. In our discussion of the spelling-pronunciation correspondence below we take this fact to indicate that the mid-14th century was a cut-off point in spelling-pronunciation contiguity, and, therefore, it gives support to the assumption that, until that point, spelling and pronunciation went largely hand in hand.

G. innovative final *-s*: 1st person, Group II, present and preterite indicative

The variation between a new, syncretic ending *-s* and nonsyncretic zero for the 1st person singular with Group II verbs, from the 14th century (Marchello-Nizia, 1992:201). This change is indeed delayed compared to the spread of *-e*.

181 Marchello-Nizia (1992:202) observed that, in the case of the stems ending in a
 182 vowel, it takes longer for the new variant to establish itself. There is also a
 183 limited number of verbs with stems ending in *-s* for etymological reasons (e.g.,
 184 *finis* < Lat. *finisco* ‘to finish’).

185 H. innovative final *-s*: 1st person, Group I, imperfect and future conditional

186 I. innovative final *-s*: 1st person, Group II, imperfect and future conditional

187 These changes can be used to model phonological changes at least until the 14th
 188 century. One of the strongest arguments in support of spelling reliability for
 189 phonological reconstruction is the novel observation, which we will discuss in
 190 more detail below, that the dropping of the final *-t* in verbs with stems ending in
 191 *u/i* is abruptly arrested and reversed just after the mid-14th century, when the
 192 French Royal Chancellerie is known (first mention 1342) to have introduced
 193 exams for the scribes requiring them to adhere to the standardized spelling rules
 194 (De Jong, 2006:25). While spelling unification had been taking place already
 195 for several decades, De Jong (2006) observed a sharp increase in what she
 196 called “parasitic consonants” after around 1340, which she attributed to the
 197 prescriptions of the official examiners. Consequently, after that point, we can
 198 only estimate verbal syncretism based on the change trajectory in the
 199 manuscripts written before that date.
 200

201 *Quantifying the emergence of the new endings*

202
 203 To establish the temporal profile of the surface changes in verbal endings, we
 204 calculated the ratio of the “new” endings to the sum of the new and “old”
 205 endings for each text in the corpus. In order to be able to identify the subject’s
 206 person in an automated way, we limited ourselves to clauses with overt nominal
 207 or pronominal subjects. This means that we took a subset of all the cases of new
 208 endings appearing in the corpus. In order to determine whether considering only
 209 overt subjects skews the results, we look at ending choice in a sample of clauses
 210 with null subjects manually annotated for subject person and conclude that there
 211 is no significant difference in the rate of new endings between null versus overt
 212 subject contexts. Thus, we can confidently estimate the rise of the new endings
 213 from a sample of clauses with overt subjects. Figures 2 and 3 show the rise of
 214 new endings divided into two major groups, namely, final *-t* deletion and *e*-
 215 insertion on the one hand, both of which resulted in an innovative *-e* ending, and
 216 *s*-insertion on the other. (Observation numbers together with a proportion of the
 217 new endings in each text are given in Tables 5B–9B, Appendix B).

218 Comparing now Figures 2 and 3 with Figure 1, on the assumption that the
 219 spelling innovations reflected changes in the verbal agreement phonology, there
 220 is no reason to assume that there was a temporal lag between the emergence of
 221 new syncretic endings and the rise of overt pronominal subjects. However, we
 222 see that, while the appearance of new *-e* and *-a* endings roughly parallels the
 223 emergence of overt subjects, innovative zero and *-s* follow a very different trend.
 224 The next question is whether we can establish a nonaccidental relation between
 225 the rise of new endings and overt subjects.




Fig. 2 - Colour online, B/W in print

 FIGURE 2. Innovative *-e* (changes A, B, C) and *-a* endings (changes D, E).

Note: $P(\text{END} = \text{NEW} \mid \text{DATE} = \text{D}, \text{GR} = \text{I}, \text{P} = 1)$ stands for the estimated probability of a Group I verb to have a new ending (i.e., *-e*) in the context of the 1st person singular subject, etc.

CLAUSE-LEVEL RELATION MODEL

We will first explore a classic line of analysis that relates null subjects and nonsyncretic agreement via a certain structural property giving rise to both; let us call it *Agr* head. The two changes are thus viewed as a consequence of the loss of the grammar with *Agr* head. We show that an approach that maintains a clause-level relation between subject expression and the type of ending makes incorrect predictions about the rise of the new endings and overt pronominal subjects. We will then suggest a more flexible approach whereby syncretic endings, rather than being a direct manifestation of an alternative structure without *Agr*, are consequences of an independent phonological change that favors the alternative grammar. Thus, the second approach dissociates null subjects from a particular set of endings in terms of surface observations, but maintains that syncretization eventually led to the disappearance of a grammar-generating null subjects.

AgrP-Grammar

As part of the first model, we assume that the initial grammar was characterized by the presence of a person feature-specified head *Agr*.⁴ We will assume that person




Fig. 3 - Colour online, B/W in print

FIGURE 3. Innovative zero (change F) and -s ending (changes G, H, I).

features introduce conditions on the denotation of a pronoun. A long semantic tradition ascribes to such features the status of presupposition triggers (e.g., Cooper, 1983; Heim, 2008; Heim & Kratzer, 1998; Kratzer, 2009; Sauerland, 2008). In addition to that, we will assume that a pronoun needs to be accompanied by an element triggering a presupposition about its reference, whether it comes as part of the morphological form of the pronoun itself or as a verbal ending.⁵ Taking the existence of the constraint for granted, we propose that person features on Agr introduce presuppositions about the subject's reference. In the absence of Agr *pro* will be left uninterpreted.

TP-Grammar

We model the replacement of null subjects with overt ones and of old endings with new ones as a passage from the initial AgrP-Grammar to an alternative grammar where verbal endings correspond to the spellout of head T, unspecified for the person feature.⁶ Since T does not carry person features, it does not introduce presuppositions necessary for a felicitous use of a *pro*.

If TP-Grammar replaces AgrP-Grammar, null subjects will become unavailable. Assuming the Constant Rate Hypothesis of Kroch (1989), our model predicts that

316 the rate of replacement of AgrP-Grammar by TP-Grammar should be the same
 317 whether it is measured as the rise of overt pronominal subjects or of new
 318 syncretic endings. For the general case, the Constant Rate Hypothesis (CRH)
 319 states that a grammatical change has the same rate of spreading in all
 320 grammatical environments, where the rate is taken to correspond to the slope
 321 coefficient of a logistic regression model. However, Kauhanen and Walkden
 322 (2017), following up on the discussion in Paolillo (2011), pointed out that the
 323 standard way of assessing statistical significance (Kroch, 1989; Pintzuk, 1995;
 324 Santorini, 1993) of a putative Constant Rate effect is statistically unsound: “if
 325 the result is not statistically significant, then it is concluded that there is support
 326 for a [Constant Rate Effect]. However, it is not sound to treat a nonsignificant
 327 value as evidence for the null hypothesis, since it was assumed to begin with.”
 328 We will maintain therefore that, whenever the result of an independence test on
 329 regression coefficients is nonsignificant, it *does not contradict* the CRH; rather it
 330 provides direct evidence for it.

331 Thus, we expect the rates of the emergence of overt pronominal subjects and of
 332 the new endings to be not significantly different. One caveat of the prediction is that
 333 even stable null subject grammars allow for overt subjects. This makes it
 334 impossible to classify a given overt pronominal subject as an instance of AgrP-
 335 Grammar or TP-Grammar, since both of them are expected to generate overt
 336 pronominal subjects. The only context that sets the two apart clearly are
 337 expletive subjects, which are consistently null in null-subject languages (e.g.,
 338 Jaeggli & Safir, 1989).⁷ We therefore will compare the rise of overt expletive
 339 subjects with the rise of the new endings. There are at least three other
 340 immediate predictions. First, the rise of the new endings should proceed at the
 341 same rate in different contexts: if the emergence of the new endings reflects the
 342 disappearance of Agr, on the CRH we do not expect this change to proceed
 343 differently depending on the verb type or the subject person. Second, there
 344 should be no increase in the frequency of null subjects in the contexts of new
 345 syncretic endings. This is so because the AgrP-Grammar that, by hypothesis, is
 346 the only grammar that can license null subjects, is associated with spellout rules
 347 which do not output syncretic endings, such as *-e* in the context of the 1st and
 348 3rd person subjects, overt or null. Finally, there should be no increase in subject
 349 expression with old, nonsyncretic endings: although AgrP-Grammar, associated
 350 with nonsyncretic endings, does sometimes generate overt subjects, their
 351 distribution is governed by constraints that produce the same rate of subject
 352 expression during the course of existence of grammar AgrP.

354 PERFORMANCE OF THE CLAUSE-LEVEL RELATION MODEL

356 *Testing the main hypothesis*

357
 358 In order to evaluate the hypothesis that the emergence of overt expletive subjects
 359 and syncretic verbal endings are two manifestations of the disappearance of the
 360 grammar with a person feature-specified Agr head, we fitted the data on the




Fig. 4 - Colour online, B/W in print

FIGURE 4. Spread of new endings and overt pronominal subjects.

appearance of overt expletives and the new endings to logistic regression models plotted in Figure 4 (parameter estimates in Table 1). The model ENDING predicts whether the verbal ending, *Y*, is *new* (or syncretic) by contrast with an *old* or nonsyncretic verbal ending as a function of time.⁸ We compare this model with an EXPLETIVE SUBJECT model that predicts whether the expletive subject realization, *Y*, is *new* (or overt) by contrast with an *old* realization where the pronominal subject is null. For the sake of comparison, we also plotted the data on the overt personal pronominal subjects.

The coefficients are not very different from each other but not identical either. To further test the CRH, we test for the contribution of the slope by comparing two mixed-effect models. The first predicts the new form *Y*, whether it is an overt expletive subject or a syncretic verbal ending, by contrast with an old form, that is, a null subject or a nonsyncretic verbal ending. The prediction is still a function of time, but we also add a random intercept α_c for each context *c*: either a morphological context or a subject context.⁹ Informally, this model means that the global model intercept may be further parametrized for each specific context, but the slope is constrained to be identical for both contexts. We compare this model to an extended version, where this time we add a random slope β_c , thus allowing the slope to vary for each context. Since the slope models the rate of

TABLE 1. *Logistic regression estimates for the new endings and overt pronominal subjects (numbers of observations of null and overt expletive and personal pronominal subjects are given in Table 4B in Appendix B)*

MODEL	COEFFICIENT	STD. ERROR	Z-VALUE	Pr(> z)	ACCURACY
ENDING	0.0062	0.0001	42.75	$< 2 \times 10^{-16}$	0.76
EXPLETIVE SUBJECT	0.0055	0.0002	25.34	$< 2 \times 10^{-16}$	0.61
PERSONAL SUBJECT	0.0063	0.0001	60.87	$< 2 \times 10^{-16}$	0.81

change, this second model allows the rate of change to differ for each context. We test whether the slope introduces a significant difference between the two models (with a log likelihood ratio test which is χ^2 distributed [$df = 2$]). The test has $p = 0.04$, and so we conclude that the introduction of the slope does better predict the data, and thus, on the CRH, these results are not compatible with the analysis of the two diachronic phenomena as stemming from the same grammatical change, which we identified as a passage from a grammar with Agr head to a grammar without it. In the remainder of this section, we will explore three other predictions made by the clause-level relation model and show that none is borne out.

Syncretization in different contexts

The model for agreement syncretization merges nine different syncretization patterns (see Appendix A). If syncretization is a consequence of the TP-Grammar associated with the new spellout rules winning over the old AgrP-Grammar, then these developments are expected to have the same rate. In order to test this, we modeled them separately, as illustrated in Figure 5 (Table 11B in Appendix B shows the estimates).

Upon visual inspection, we see that the spread of the new ending *-e* has more or less the same profile in all of its contexts. In contrast, it differs from the spread of *-a* and *-s*, contrary to what was predicted by the clause-level relation model. Thus, individual endings spread at different rates, and the innovations seem to group into classes in terms of their phonological environments.¹⁰

Spread of the new endings with null subjects

Another prediction made by the clause-level relation model is that there should be no increase in the new endings in the context of null subjects. We do find occurrences of *-e* in the context of the 1st or 3rd person singular null subjects (see Table 9B in Appendix B), yet such occurrences of new endings with null subjects are not frequent: at all times they stay below 20 per text. One way to explain away their occurrence is to analyze them as etymological vowels that create noise in the passage from the old to new endings. However, if that is indeed noise, we expect it not to become stronger with time. To test this




Fig. 5 - Colour online, B/W in print

FIGURE 5. Logistic regression models of the emergence of the new endings.

expectation, we fit the data on the appearance of *-e* in the context of the 1st person singular overt and null subjects to a logistic regression. As Figure 6 shows, the trend is the same (see Table 12B in Appendix B for the estimates).

This result is unexpected if *-e* with null subjects is just an etymological residue. Rather, the observation that the new ending spreads at similar rates in the context of null and overt subjects suggests that we are witnessing one and the same (phonological) change operating in different contexts. In other words, the choice of ending is independent of the expression of the subject, contrary to what is predicted by the structural model relating subject expression and ending type as manifestations of a particular grammar. Note that we do not need to check the spread of different types of new endings with overt and null subjects, since the clause-level relation model predicts that no new endings increase with null subjects and is therefore falsified even by one case of the contrary.

Spread of overt subjects with old endings

The final prediction that we derive from the clause-level relation model is that there should be no increase in subject expression in the context of verbs with old, nonsyncretic endings. We compared the rate of subject expression in the




Fig. 6 - Colour online, B/W in print

FIGURE 6. Rise of *-e* with Group I verbs in the context of the 1st person singular subjects.

contexts of verbs with the old nonsyncretic endings *-t*, zero on the one hand and new syncretic endings *-e*, *-s* on the other. We estimated the probability of having an overt pronominal subject for finite clauses with verbs ending in *-e* (Group I & II), *-t* (Group I & II), *-s* (Group II), and zero (Group I & II) endings, as shown in Figure 7 (Table 13B in Appendix B).¹¹ Clearly, the subject expression rate grows over time for the nonambiguous endings.¹² Relatedly, Ranson (2009) concluded, based on the three texts she examined, that ending ambiguity is not a good predictor of subject expression.

In sum, we have shown that a number of predictions generated by a model that assumes that subject expression and agreement type are related at the clause level via a certain functional head are not borne out. Namely, new endings spread at different rates depending on the ending type, which is unexpected if both are generated by a new grammar that is supposed to spread at the same rate on the CRH. In addition, new endings spread both with overt and null subjects, contrary to the model's assumption that null subjects are generated only by the old AgrP-Grammar, where the Agr head spells out as old, nonsyncretic endings. Finally, the expectation that there would be no increase in overt subjects in the context of old, nonsyncretic endings, which, by hypothesis, are generated by the AgrP-Grammar producing overt subjects at a constant (relatively low) rate, is




FIGURE 7. Pronominal subject expression with old and new endings.

also not borne out. The overall conclusion is that a model that assumes a strict dependency at the clause level between what type of endings are used and whether or not pronominal subject is expressed is not supported by the diachronic data. However, we need to deal with another possible explanation for why we do not find a complete parallelism between ending syncretization and pro-drop disappearance, namely, that the verb ending changes registered in the written texts are not reflective of the phonological reality and therefore cannot be used to evaluate a clause-level relation hypothesis.

SPELLING-PRONUNCIATION PROBLEM

For the purposes of the present study, the problem of the correspondence between pronunciation and spelling entails two independent questions. The first one is whether the spelling innovations had phonological substance. The second question about the spelling-pronunciation relation is concerned with the emergence of phonological innovations behind conservative orthography. The state of Modern French witnesses the fall of all the stops and sibilants (at least in an isolated pronunciation) that used to correspond to the present-day word-final

Fig. 7 - Colour online, B/W in print

541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585

586 consonantal graphemes, not just the final *-t* whose disappearance we tracked above.
 587 Again, judging from the Modern French spelling-pronunciation correspondence,
 588 this change is mostly not reflected in spelling. For the second part of our study,
 589 where we attempt to estimate the general level of syncretism in the system, it is
 590 important to know until what point in time we can equate presence in the
 591 spelling with phonological presence. Fortunately, it seems that we can estimate
 592 this date with a great deal of precision due to the co-occurrence of two
 593 independently attested facts. First, there exists a historical record of the first
 594 centralized spelling standardization in the mid-14th century. Second, our data
 595 show that the disappearance of the final *-t* with Group II verbs with unstressed
 596 roots ending in *-i/-u*, which, if it had followed a statistically expected trajectory,
 597 would have reached its completion around that time, was stopped and reversed
 598 in the late 14th century (Figure 8). This presumably shows the effect of the
 599 spelling standardization that marked the end of the strict spelling-pronunciation
 600 correspondence.

601 There seems to be a consensus that the rise of the new endings reflected the
 602 phonological reality rather than simply a change in orthographical conventions
 603 (Dees et al., 1980; Fouché, 1931; Goyette, 1993; Marchello-Nizia, 1992; Morin,
 604 2001; van Reenen & Schøsler, 1987). Arguments against a possible claim that
 605 what we observe in texts is just variation in writing conventions can be divided
 606 into the following groups. First, as we saw, the emergence of *-e* as a final
 607 grapheme in the context of Group I verbs with a 3rd person singular subject in
 608 present indicative and subjunctive follows a logistic curve whose slope is
 609 indistinguishable from the slope of the curve modeling the emergence of *-e* with
 610 Group I verbs in the context of 1st person singular subjects in present indicative
 611 and subjunctive.¹³ These results fit well with the hypothesis of Dees et al.
 612 (1980) and van Reenen and Schøsler (1987) about /e/-insertion being a
 613 compensatory process triggered by the instability of the final stops to preserve
 614 the integrity of the root. The appearance of *-e* as a final grapheme with the 3rd
 615 person singular subjects on this view results from the fall of the final /t/ (e.g.,
 616 *aimet* > *aime* ‘(he) loves’ and *aint* > *aime* ‘(he) would love’), whereas its
 617 appearance with the 1st person singular subjects results from a compensatory
 618 /e/-insertion to keep the root final consonants from not being pronounced (e.g.,
 619 *aim* > *aime* ‘(I) love’ and ‘(I) would love’). Although the quasi-identity of
 620 slopes is only indicative, this is expected on the hypothesis that this is a
 621 paradigm-wide morphophonological process. That is, given the CRH, it is
 622 entirely expected for a morphophonologically conditioned change to proceed at
 623 the same rate in different environments (cf., Fruehwald, Gress-Wright, &
 624 Wallenburg, 2009). Second, according to our estimates, in the context of the 3rd
 625 person subjects syncretization happened earlier than with the 1st person singular
 626 subjects, which makes sense if the fall of the final stops that were not part of the
 627 root (again, *aimet* > *aime* ‘(he) loves’ and *aint* > *aime* ‘(he) would love’)
 628 preceded the emergence of a “compensatory” /e/ following root-final
 629 consonants. In contrast, on the hypothesis that what we observe are changes in
 630 spelling conventions, although not theoretically impossible, it would look like a




Fig. 8 - Colour online, B/W in print

FIGURE 8. Change reversal for Group II verbs in preterite with 3rd person subject.

series of strange coincidences if, first, different spelling conventions in different contexts were changing at very same rates, and, second, if they first changed in the context of the 1st and then of the 3rd person subjects. Third, according to Fouché (1931:180) and Marchello-Nizia (1992:201), in the context of the 1st person singular subjects in indicative and subjunctive, the *-e* grapheme first spread in the context of consonant-final and only later vowel-final roots (e.g., *cri-er* ‘to shout’ where in the context of the 1st person singular subjects *cri* was replaced by *crie*). Again, this fits a phonology-based account of the change, since the sequence of spreading across contexts can be described in terms of phonologically natural classes, whereas this appears as a mysterious plotting of the scribes on the spelling convention-based account. Lastly, a phonologically motivated change affecting final vowels has precedence in the history of Late Latin, where all the final vowels ended up falling except for those cases where their fall would have led to an unacceptable consonant cluster (see the discussion in Goyette [1993] and references therein). In Old French, reflexes of this process are the so-called “etymological e,” that is, root-final *-e* following certain consonant clusters, as in *siffle-r* (from Latin *sibila-re* > *sifila-re* > *sifla-re*) (cf., *don-er* from Latin *dona-re*, which lost its root final /a/ in Late Latin, unlike *siflare*). It is not so surprising then to see another round of “compensatory” root

676 final /e/, this time as an epenthetic process meant to keep the root final consonants
 677 from falling. In view of these arguments, none of which supports an account of the
 678 new endings in terms of spelling convention changes, we conclude that our results
 679 based on written source can be plausibly projected onto the phonological reality
 680 and thus used to test a model structurally relating syncretism, introducing changes
 681 and pro-drop disappearance. Similarly, Fruehwald et al. (2009) analyzed data on
 682 the loss of final fortition in (Bavarian) Early New High German, observable in
 683 orthographic variation of the period, for example, *tak* versus *tag* ‘day (acc. sg),’ *rat*
 684 versus *rad* ‘counsel (acc.sg),’ and argued that this variation clearly represents a
 685 phonological change in progress rather than shifting scribal tradition. When it
 686 comes to determining at what point the phonological reality behind conservative
 687 spellings changed, the reconstructions of the timing of the fall of the final
 688 consonants rely mostly on the analysis of rhymes (matching versus nonmatching),
 689 hypercorrections (insertions of etymologically absent consonants), omissions of
 690 etymologically present consonants, commentaries in the grammars of the time, and
 691 analyses of the borrowing from French into other languages that likely reflected the
 692 spelling at the time of borrowing. The dating question is important to us in as
 693 much as we want to take into account final consonant instability when evaluating
 694 the overall degree of syncretism or ambiguity in the verbal system. De Jong (2006)
 695 undertook a statistical analysis of the rhymes in three texts written in the Parisian
 696 dialect in the 13th-14th centuries. She looked at the frequency of the nonmatching
 697 rhymes for a given grapheme (e.g., *escript* ‘text’-(*je*) *pris* ‘I take’) compared with
 698 that of the matching rhymes (e.g., *moult* ‘many’-(*je*) *doubt* ‘I doubt’), taking
 699 higher-than-chance frequencies to be indicative of the grapheme nonpronunciation.
 700 One of the general conclusions of De Jong (2006:176) is that the nonpronunciation
 701 of the final consonantal graphemes increases dramatically in the 14th century. This
 702 is the period when the mismatching rhymes, including mismatches involving our
 703 consonants of interest, begin to be observed in her corpus (cf., Foulet, 1935).
 704 Importantly, De Jong (2006:174) linked the emerging mismatch between spelling
 705 and pronunciation with a particular historical event, namely, the introduction by the
 706 Royal Chancellery in Paris of the standard exams for the scribes in 1342. We
 707 found a rather dramatic argument in favor of this hypothesis in the form of the
 708 reversal of the final -*t* disappearance in the preterite forms of certain Group II verbs
 709 (Figure 8). We cannot conceive of any plausible explanation of this development in
 710 phonological terms. Rather, it seems to result precisely from an artificially
 711 introduced norm.

712 CHANGE AS A VARIATIONAL LEARNING OUTCOME

713
 714
 715
 716 As was demonstrated above, the long-standing intuition going back to at least Foulet
 717 (1928) that it was the impoverishment of the verbal endings that triggered the loss of
 718 null subjects cannot be implemented as a model in which non-syncretic endings and
 719 null subjects are considered manifestations of the same grammatical property.
 720 However, given that overall the new endings and overt pronominal subjects

(whether personal or expletive) spread at almost the same rates, illustrated in [Figure 4](#), it would likewise be counterintuitive to conclude that we should give up altogether on all the models which assume a non-accidental relation between the two changes.

Sprouse and Vance (1999) proposed the first, to our knowledge, reinforcement learning model to explain the loss of null subjects, appealing to the processing difficulty associated with their parsing. In this model, a null subject has a greater chance of inducing a parsing failure than its competitor, an overt pronominal subject. Since, by the authors' assumption, speakers tend to produce grammatical forms at frequencies at which they have encountered them in their speech community, failures to parse null subjects will lead to the decrease in the frequency of null subjects in the output of the speakers, which in turn will reduce the ambient frequency of null subjects on the next cycle. The cycle repeats until null subjects vanish from the speech community.

Below we suggest a model of the loss of null subjects which builds on the variational learning model proposed in Yang (2002, 2010). Ambiguous endings are considered within this model as the main factor that creates a parsing difficulty for null subjects (contra Sprouse & Vance [1999]).

General framework

Yang's (2002, 2010) model is based on the assumption that children have innate access to multiple grammatical systems and, in the course of language learning, use the input data to probabilistically evaluate the available options. They may either converge on a single grammar, or, as adults, they may end up with multiple grammars used at certain probabilities, which corresponds to the case of synchronic variation. Depending on whether the next generation arrives at the same or different probability distribution, we get the case of diachronically stable variation or diachronic change respectively. Hypothesizing what kind of data contributes to the probabilistic evaluation of the grammars, we can approximate the course of the competition based on corpus distributions of the relevant data.

Formally, we use Yang's (2002, 2010) model as a way to estimate the probabilities $P(\mathbf{G} = G_1)$ of using the grammar G_1 and $P(\mathbf{G} = G_2)$ of using the grammar G_2 from a data set $X = x_1 \dots x_n$ in which, for a specific example $x \in X$, we are not sure which of G_1 or G_2 actually generated x . Informally, the estimation procedure is iterative and increases $P(\mathbf{G} = G_i)$ when G_i successfully parses an example x while it decreases $P(\mathbf{G} = G_j)$ ($i \neq j$). The iterative procedure runs as follows:

- Select randomly a clause x in the data set X
- Select randomly G_i in proportion to its probability
- Analyze x with G_i
 - If G_i succeeds in analyzing x , provide G_i a reward and G_j a penalty: $P(\mathbf{G} = G_i)$ increases and $P(\mathbf{G} = G_j)$ decreases.
 - If G_i fails in analyzing x , provide G_i a penalty and G_j a reward: $P(\mathbf{G} = G_i)$ decreases and $P(\mathbf{G} = G_j)$ increases.

Using the notation $G_i \not\Rightarrow x$ to indicate that G_i fails to parse x , we can define the notion of penalty of a grammar G_i as $c_i = P(G_i \not\Rightarrow x)$. That is, c_i is the probability that G_i fails to analyze an example in X . This quantity can be estimated simply by counting the proportion of a grammar's failures in the data set. Given this notion, for the case where we have two grammars G_1, G_2 with penalties c_1, c_2 , Narendra and Thathachar (1989) proved the following theorem:

$$\lim_{t \rightarrow \infty} P(G = G_1 | T = t) = \frac{c_2}{c_1 + c_2} ; \lim_{t \rightarrow \infty} P(G = G_2 | T = t) = \frac{c_1}{c_1 + c_2}$$

The probability of using a grammar G_i is proportional to the number of observed failures of G_j in the data set ($i \neq j$). Specifically $P(G = G_i) = 1$ when G_j always fails and $P(G = G_i) = 0$ when G_j never fails.

Diachronic stability and change

The outcome of the learning process (possibly over the lifespan) may stay the same or it may change from one generation to another. In the model we are considering, the only reason why learning may not converge on grammar G_i is if its penalty probability c_i is greater than zero, that is, if there are some subset input data that G_i fails to parse. Once c_i associated with G_i becomes greater than zero, a language may leave a diachronically stable state and enter a state of diachronic change. Moreover, an increase in the frequency of the data unparseable with G_i in the next generation will lead to the increase in c_i , and so on to the point when G_i gets completely demoted. Emergence of such data may have nothing to do with the grammatical options themselves and may stem from phonological changes as well as from a second language interference.

Applying this to the loss of null subjects in Medieval French, let us assume that the initial winning grammar (Agr-P Grammar) is the one that licenses null pronominal subjects. Its competitor (the TP-Grammar) only generates clauses with an overt subject. Notice that this model incorporates the Taraldsen/Rizzi insight about a categorical, core grammar-based dependency between functional head features and null subjects. In order to model the competition between these two grammars, the crucial parameters are the penalty probabilities of the grammars. By hypothesis, AgrP-Grammar fails each time the information about a subject's reference cannot be retrieved from the verbal ending, which is the case whenever the ending is ambiguous. An ending is classified as ambiguous in case the speaker has been exposed to a data sample where the ending occurs in the context of overt subjects with various (more than one) person specifications.

In the case of ambiguous endings, the Agr head cannot be projected during the parse, since there is not enough information to give it semantic content. In contrast, TP-Grammar fares well with all kinds of endings (as long as tense information can be read off of them), but fails when chosen to parse null-subject clauses. In those cases in the absence of a subject DP providing presupposition triggering features, the domain of the external argument of the verb is left underspecified, and the composition does not converge. Now a diachronically stable null subject

situation is predicted to obtain in case there are no problematic data of the kind described above, that is, there are no ambiguous endings and the penalty probability c_{Agr} is 0. This means that AgrP-Grammar never fails and in every generation ends up driving the competing TP-Grammar out, since the latter cannot parse some of the AgrP-Grammar's output, namely null-subject clauses.

Estimating failure probabilities

To estimate c_{Agr} , we exhaustively classify verbal endings as ambiguous or unambiguous. We define an ending as ambiguous if it does not correspond to a unique combination of person and number features (see Appendix A). We coded every finite clause in the corpus (as usual, with the exclusion of subject *wh*-clauses and imperatives) as to whether the verbal ending is unambiguous. In the case of endings that were ambiguous already in the earliest texts, all clauses with a finite verb having such an ending have been coded as ambiguous. In the case of endings classified as having emerging ambiguity, namely, those that became ambiguous later than in the earliest texts, we classified clauses dated before the first attested cases of ambiguity as having unambiguous predicates and those dated after the ambiguity emerged as having ambiguous predicates. The failure probability c_{Agr} is then estimated as the frequency of clauses with ambiguous predicates at a given date. In the case of TP-Grammar the estimate of c_{TP} is even more straightforward: it is the frequency of null-subject clauses.

The predicted value of PTP given c_{Agr} and c_{TP} as estimated using the matrix above is plotted in [Figure 9](#) (observation numbers on which c_{TP} and c_{Agr} are based are given in Appendix B, Tables 4B and 10B, respectively).

Discussion

We estimated the parsing probability of the TP-Grammar based on our estimates of the probabilities with which this grammar and its competitor, AgrP-Grammar, encounter data that they cannot parse. The parsing probability of the TP-Grammar grows steadily during the Medieval period.

Recall that the parsing probability of the TP-Grammar in the limit is the ratio of the probability of AgrP-Grammar to fail to the sum of the AgrP-Grammar and TP-Grammar probabilities to fail. This means that the greater the probability of the AgrP-Grammar to fail, the greater, eventually, will be the probability of the TP-Grammar to be used. Given how we estimate the AgrP-Grammar's probability to fail, that is, as the frequency of ambiguous endings, it is clear that our estimate of the TP-Grammar probability to be used is dependent on the frequency of ambiguous endings. Thus, this model, without assuming that ending ambiguity and subject expression depend on the same underlying factor, puts the two in a relation of direct dependency. This is a welcome configuration given the desiderata expressed above, namely, finding a model which would dissociate the two phenomena at the clause level but would relate them in the course of language evolution. It is worth stressing here that a given ending in this model does not reveal which grammar was used to generate it: by assumption, the




FIGURE 9. Parsing probability of the TP-Grammar.

spread of syncretic endings is a phonological phenomenon which is “blind” to the syntactic origins of the string it is operating on.

As Figure 10 shows, the curve corresponding to the parsing probability of the TP-Grammar is roughly parallel to the estimated probabilities of personal and pronominal subject expression. Intuitively, the probability of the expletive subject expression corresponds to the probability of the TP-Grammar to be chosen for production, since this is the only grammar that can generate overt expletive subjects. What we observe, then, is the TP-Grammar parsing probability lagging significantly behind its production probability. One explanation for the apparent lag is that our estimate of the ending ambiguity, to which the TP-Grammar production probability is directly related, is overly conservative. We have already mentioned on several occasions that the spelling standardization of the mid-14th century seems to have had a visible effect on the manuscripts upon which our corpus is based. The phenomenon we discussed is the reintroduction of the final *-t* for the Group II verbs with unstressed roots in *-i/-u*. Given the disappearance of the inflectional *-t* after vowels prior to 1300, the same should have happened to the *-t* following glides and stops shortly thereafter (cf., a suggestion in Buridant [2000:250]). And we know for a fact that eventually final inflectional stops did fall in all the environments. However,




Fig. 10 - Colour online, B/W in print

FIGURE 10. Parsing probability of the TP-Grammar (non-conservative*).

*Note: Assuming that, in addition to *-e*, *-s*, zero following *i/u*, and *-eies*; also ambiguous were *-t*, *-eiet*, *-it*, *-et*, *-at*.

presumably due to the spelling standardization, we do not observe these changes and, therefore, cannot take them into account in our estimates of ending ambiguity, which gives the impression that the latter lags seriously behind the production probability reflected in the rate of overt expletive subjects.

We can see what happens to the parsing probability of the TP-Grammar if we make a less conservative assumption about final consonant fall. That is, let us assume that, in addition to the endings *-e*, *-s*, zero following *i/u*, and *-eies*, the following endings were ambiguous as well by virtue of effectively not being pronounced from 1400 on and thus resulting in verbal forms homophonous with either 1st or 2nd person singular forms: *-t*, *-eiet*, *-it*, *-et*, *-at*. In Figure 10, one can see that this less conservative estimate is almost identical to the estimated production probability in the form of overt expletives.

One may argue, however, that this parallelism, in general, is not a particularly interesting result, since, in addition to the frequency of ambiguous endings, the TP-Grammar parsing probability in the limit inversely depends on the relative frequency of null subjects, which, of course, decreases over time. Consequently,

the question is whether ending ambiguity actually plays an important role in predicting the TP-Grammar parsing probability.

One way to evaluate the role of the endings for the outcome of the grammar competition is to design a variational learning-based model in a way that would not make reference to them at all and to compare this “ending-less” model to the model that does take them into account, such as the one that we have just considered. To this end, we use the measure of grammar fitness proposed in Yang (2000) and used, in particular, in variational learning models of the loss of V-to-T raising in Scandinavian in Heycock and Wallenberg (2013) and the loss of OV in Latin in Danckaert (2017). Fitness of grammar G is defined as the proportion of clauses that only G generates out of all clauses that G generates, or the proportion of unambiguous clauses in the output of G .

Fitness of the AgrP-Grammar cannot be straightforwardly estimated from our data, since, by hypothesis, all the attested stages of historical French correspond to mixed grammar states, that is, to the outputs of the two competing grammars. This follows from the assumption that, whenever we find overt expletive subjects, the TP-Grammar must have been at work and from the fact that expletive subjects are found in the earliest attested texts (e.g., Prévost, 2018; Zimmermann, 2014). Instead, we can approximate the fitness measure of the AgrP-Grammar on the basis of a language that is currently in a “pure” pro-drop state. The estimated probability of null subjects is around 0.7 for pro-drop languages such as Italian or Spanish (e.g., Bates, 1976; Nagy, Aghdasi, Denis, & Motut, 2011; Otheguy, Zentella, & Livert, 2007:778). That is, by assumption, the AgrP-Grammar produces unambiguous clauses with the probability of 0.7. Now the fitness of TP-Grammar corresponds to the estimated probability of expletive subjects in a nonpro-drop language, such as English, which obviously cannot be anywhere near 0.7.¹⁴ Given these approximations, $\text{Fitness}(\text{AgrP-Grammar}) \gg \text{Fitness}(\text{TP-Grammar})$. By the Fundamental Theorem of Language Change (Yang, 2000:239), which states that the winner in the long run is always the grammar with a greater fitness, this model predicts that AgrP-Grammar wins hands down, contrary to the historical facts. We thus conclude that a model that factors in ending ambiguity fares better than a model that does not, which supports the assumption that there is a causal relation between ending syncretization and null-subject disappearance.

CONCLUSIONS

The goal of this paper was to bring parsed corpus data and statistical modeling to bear on the old-standing puzzle of the relation between the disappearance of null subjects and verbal subject agreement syncretization in the historical development of French. We engaged the Constant Rate Hypothesis in order to explore a model that relates the two changes as reflexes of one underlying structural shift and showed that it generates predictions that are not supported by the data. The key feature of the failed predictions is the independence of the two

991 developments. Specifically, we found that the increase in overt personal
 992 pronominal subjects was uniform across old *and* new endings, and, likewise, that
 993 the spread of the new syncretic endings was uniform across clauses with null
 994 *and* overt subjects.

995 A second model that we explored related the two changes via the step of
 996 language learning whereby one change (ending syncretization) promotes the
 997 appearance of sentences that disadvantage a grammar with a null-subject option
 998 and thus automatically favor the overt-subject grammar, thus causing null-
 999 subject disappearance. We approached this model from two perspectives.

1000 First, we focused on what we assumed to be the parsing capacities of the two
 1001 competing grammars and used the linear reward-penalty theorem to estimate the
 1002 evolution of the probabilities of the two grammars over time by estimating their
 1003 failure probabilities at each time point. Crucially, the failure probability of a
 1004 null-subject grammar is taken to be directly related to the frequency of
 1005 ambiguous endings in the data. This was the core assumption meant to capture
 1006 the intuitive link between agreement quality and subject expression in the
 1007 process of language change. Another crucial assumption, which made the model
 1008 compatible with the facts concerning the surface-level independence of subject
 1009 and ending types, was that the type of ending is determined by a phonological
 1010 process that is entirely independent of which grammar is picked by the speaker
 1011 to generate a given sentence.

1012 As a second possibility, we estimated fitness of null and overt subject grammars
 1013 based on the data from pure state languages. This time the measure did not rely on
 1014 ending ambiguity. Estimated this way, fitness gives an advantage to the null-subject
 1015 grammar, and on Yang's Fundamental Theorem it is expected to win, contrary to
 1016 the historical facts. We conclude, thus, that so far the best model of null-subject
 1017 disappearance is one that factors in the increase in ending ambiguity without
 1018 assuming a categorical clause-level dependency between the two phenomena.

1019 Synchronic studies of variation in null-subject expression in Romance
 1020 languages, to the best of our knowledge, fail to establish verbal ending
 1021 ambiguity as a relevant factor, and thus leave us with a puzzle as to the nature of
 1022 Taraldsen's generalization. For instance, Nagy and Heap (1998) reported that
 1023 whether an ending is ambiguous is not a good predictor of subject expression in
 1024 Francoprovençal. The same conclusion is reached in Carvalho and Child (2011)
 1025 based on Spanish material. This agrees with Ranson's (2009) conclusions and
 1026 our own observations concerning the diachronic French data. Our work thus
 1027 supplements synchronic variationist studies in that we offer a diachronic model
 1028 that can capture the relation without postulating a clause-level dependency. This
 1029 suggests that, in some cases, the study of natural language variation *must* include
 1030 the temporal dimension, otherwise some potentially highly relevant factors will
 1031 remain "invisible."

1032 It has to be noted that our conclusions do not rule out in principle a clause-level
 1033 dependency between surface forms. Such an outcome could, for instance, be the
 1034 result of a competition between Agr- and TP-Grammars, whereby by the end of
 1035 a variational learning cycle they end up in a complementary distribution with

respect to tense/aspect environments. An analysis along these lines would need to be worked out for the systems where the only contexts disfavoring subject omission are certain tense/aspectual forms syncretic with respect to subject person, such as, for instance, some Northern Italian, Franco-Provençal, and Occitan dialects (Manzini & Savoia, 2005), Hebrew (Shlonsky, 2009), Finnish (Koenenman, 2006), Irish (Speas, 1995), Russian (Bizzarri, 2015).

This study is part of a more general agenda of using diachronic material for the study of interfaces, that is, formal relations between syntax, morphology, phonology, and semantics/pragmatics. Another prominent group of what seems to be parallel and potentially related changes is the disappearance of nominal case marking and word order changes. Simonenko, Crabbé, & Prévost (2015) showed that the remnants of the case opposition in Medieval French disappear within approximately the same timeframe as the possibility of having an OV order. It remains to be seen in further research if any of the models explored above can be used to explore the nature of the relation between these two changes.

SUPPLEMENTAL MATERIALS

Appendices A, B, and C can be found at: <https://doi.org/10.1017/S0954394519000188>

NOTES

1. In addition to demonstratives, we counted as nominal subjects all subject phrases headed by a noun, both animate and inanimate, as well as nominalized adjectives, numerals, quantifiers, and free relatives, disregarding their syntactic position with respect to the finite verb.

2. There are a handful of other (less frequent) conjunctive adverbs, such as *puis*, as well as a disjunction *mais* that seem to license subject ellipsis in Modern French and that we did not exclude from our dataset, since this would require an exhaustive study of ellipsis-licensing conditions in Medieval French.

3. Here we are setting aside the question of whether subject clitics in Modern French function as subject-agreement markers (see De Cat [2005] for a discussion).

4. Importantly, we are not assuming that we can necessarily observe in existing sources the stage of French where the only grammar in use was the one we call here initial, that is, the null subject grammar without ending syncretism. This stage was likely left undocumented. This is why, we think, Old French is sometimes called a partial or nonpro-drop language (e.g., Kaiser, 2009; Zimmermann, 2009). In terms of the grammar competition model of language change that we assume, the data which only partially conform to the criteria of a given grammatical type are modeled as a mix of outputs of two “pure” grammatical types.

5. In this relation, we can invoke a long tradition, going back to Benveniste and supported by typological observations, of ascribing a universal status to the person distinctions in natural languages (e.g., Harley & Ritter, 2002; Tvica, 2017).

6. We assume that *-ez* and *-ons* endings, distinct from the rest and each other, are exponents of the feature [plural] in the context of 1st and 2nd person plural subjects.

7. As mentioned above, the stage where only AgrP-Grammar is operative is not attested; we find overt expletives in the earliest French documents, even though at a very low rate.

8. The ENDING model has the form $P(Y = \text{new} \mid \text{Date} = d) = \frac{e^{\alpha + \beta d}}{1 + e^{\alpha + \beta d}}$, where α is the intercept and β the slope. The intercept is interpreted as an abstract indicator of when the change takes place in time, and the slope is interpreted as the rate of change. The EXPLETIVE SUBJECT model has the same form, but this time Y represents expletive subject realization instead of verbal syncretism. In order to illustrate both models we first fitted them separately to the data.

9. The two models are $P(Y = \text{new} \mid \text{Date} = d, \text{Context} = c) = \frac{e^{\alpha + \beta d}}{1 + e^{\alpha + \beta d}}$ and $P(Y = \text{new} \mid \text{Date} = d, \text{Context} = c) = \frac{e^{\alpha + \beta d + \beta c}}{1 + e^{\alpha + \beta d + \beta c}}$.

1081 **10.** One way to interpret these findings is to assume that first there is a fall of the ending-final stops after
 1082 *-e-* and a compensatory /e/-insertion following root-final stops, which appears as an innovative *-e*
 1083 ending. Second, there is a fall of the ending-final stops after *-a-*, which appears as an innovative *-a*
 1084 ending. Third, there is a fall of the ending-final stops after *-i/-u* (Group II 3rd person preterite) and
 1085 /s/-insertion after *-i/-u* (Group II 1st person present and preterite), which results in innovative zero
 1086 and *-s* endings respectively. The latter process can arguably be considered as compensatory in
 1087 relation to the former in order to keep the 1/3 person distinction. Finally, there is an innovative /s/
 1088 insertion after diphthongs (Group I and II first person imperfect and future conditional).

1087 **11.** We fit the seven datasets to the logistic regression model $P(\text{Pron Sbj} = \text{yes} \mid \text{Date} = d) = \frac{e^{\alpha + \beta \cdot d}}{1 + e^{\alpha + \beta \cdot d}}$.

1088 **12.** One can see that overt pronominal subjects spread at very similar rates with all endings except for *-t*
 1089 with Group I verbs. A plausible explanation is that *-t* with Group I verbs virtually disappears after 1200,
 1090 and what we see past that date is essentially noise which skews the model.

1090 **13.** We used the same approach as for the comparison between the slopes of expletive subject and new
 1091 ending emergence models, with $p = 0.8$ this time.

1092 **14.** According to Chocholoušová (2009), in sentence-initial position in English texts, dummy subjects
 1093 occur at the frequency of 0.25% per 10,000 words. The author did not give a frequency in terms of
 1094 sentences, but if we roughly estimate an average English sentence as consisting of 20 words, this
 1095 gives us a frequency of 4.8%. That is, 0.05 can be used as a (very rough) approximation of the
 1096 probability of the unambiguous output by the TP-Grammar.

1097 REFERENCES

- 1098 Adams, Marianne. (1987). From Old French to the theory of pro-drop. *Natural Language & Linguistic*
 1099 *Theory* 5:1–32.
- 1100 Bates, Elizabeth. (1976). *Language and context: The acquisition of pragmatics*, volume 13. New York:
 1101 Academic Press.
- 1102 Bettens, Olivier. (2015). Chantez-vous français ? Remarques curieuses sur le français chanté du Moyen
 1103 Âge à la période baroque. URL <http://virga.org/cvf/>.
- 1104 Bizzarri, Camilla. (2015). Russian as a partial pro-drop language: Data and analysis from a new study.
 1105 *Annali di Ca' Foscari. Serie occidentale* 49.
- 1106 Buridant, Claude. (2000). *Nouvelle grammaire de l'ancien français*. Paris: Sedes.
- 1107 Carvalho, Ana M., & Child, Michael. (2011). Subject pronoun expression in a variety of Spanish in
 1108 contact with Portuguese. In *Selected Proceedings of the 5th Workshop on Spanish Sociolinguistics*,
 1109 14–25.
- 1110 Chocholoušová, Bohumila. (2009). *Dummy subjects in English, Norwegian and German. A parallel*
 1111 *corpus study*. Doctoral Dissertation, Masarykova univerzita.
- 1112 Cooper, Robin. (1983). *Quantification and syntactic theory*. Dordrecht: Reidel.
- 1113 Danckaert, Lieven. (2017). The loss of Latin OV: Steps towards an analysis. In *Elements of comparative*
 1114 *syntax: Theory and description* 127:401–44. Berlin: Walter de Gruyter GmbH & Co KG.
- 1115 De Cat, Cécile. (2005). French subject clitics are not agreement markers. *Lingua* 115:1195–219.
- 1116 De Jong, Thera. (2006). *La prononciation des consonnes dans le français de Paris aux 13^{ème} et 14^{ème}*
 1117 *siècles*. Utrecht: LOT/Netherlands Graduate School of Linguistics.
- 1118 Dees, Anthonij, Meilink, Steintje, van Reenen-Stein, Karin, & van Reenen, Pieter. (1980). Un cas
 1119 d'analogie: l'introduction de *-e* à la première personne du singulier de l'indicatif présent des
 1120 verbes en *-er* en ancien français. *Rapport/Het Franse Boek* 50:105–10.
- 1121 Ewert, Alfred. (1943). *The French Language*. London: Faber & Faber.
- 1122 Fontaine, Carmen. (1985). *Application de méthodes quantitatives en diachronie: L'inversion du sujet en*
 1123 *français*. Master's thesis. Université du Québec à Montréal.
- 1124 Fouché, Pierre. (1931). *Le Verbe français: étude morphologique*. Paris: Société d'édition "Les Belles
 1125 Lettres."
- 1126 Foulet, Lucien. (1928). *Petite syntaxe de l'ancien français*. Paris: Champion, troisième édition revue.
 1127 Réédition 1982.
- 1128 _____. (1935). L'extension de la forme oblique du pronom personnel en ancien français. *Romania*
 1129 61–62, 257–315.
- 1130 Franzén, Torsten. (1939). *Étude sur la syntaxe des pronoms personnels sujets en ancien français*.
 1131 Doctoral Dissertation. Uppsala.
- 1132 Fruehwald, Josef, Gress-Wright, Jonathan, & Joel Wallenberg. (2009). Phonological rule change: The
 1133 constant rate effect. In S. Kan, C. Moore-Cantwell, & R. Staubs (Eds.), *Proceedings of the 40th*
 1134 *Annual Meeting of the North East Linguistic Society*, volume 1, 219–30. GLSA Publications.

- Goyette, Stéphane. (1993). *Le système verbal de Jacques Peletier du Mans, XVIe siècle*. Master's thesis. Université de Montréal.
- Harley, Heidi, & Ritter, Elizabeth. (2002). Person and number in pronouns: A feature-geometric analysis. *Language* 78:482–526.
- Haspelmath, Martin. (1999). Optimality and diachronic adaptation. *Zeitschrift für Sprachwissenschaft* 18:180–205.
- Heim, Irene. (2008). Features on bound pronouns. In D. Harbour, D. Adger, & S. Bejar (Eds.), *Phi Theory: Phi-features across Modules and Interfaces*, 35–52. Oxford: Oxford University Press.
- Heim, Irene, & Kratzer, Angelika. (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
- Heycock, Caroline, & Wallenberg, Joel. (2013). How variational acquisition drives syntactic change. *The Journal of Comparative Germanic Linguistics* 16:127–57.
- Hirschbühler, Paul. (1992). L'omission du sujet dans les subordonnées V1: les CNN de Vigneulles et les CNN anonymes. *Travaux de linguistique* 24:25–46.
- Hirschbühler, Paul, & Junker, Marie-Odile. (1988). Remarques sur les sujets nuls en subordonnées en ancien et en moyen français. *Revue québécoise de linguistique théorique et appliquée* 7:63–84.
- Ingason, Anton Karl. (2016). PaCQL: A new type of treebank search for the digital humanities. *Italian Journal of Computational Linguistics* 2:51–66.
- Jaeggli, Osvaldo & Safir, Kenneth J. (1989). The null subject parameter and parametric theory. In O. Jaeggli and K. J. Safir (Eds.), *The null subject parameter*, 1–44. Dordrecht: Springer.
- Kaiser, Georg A. (2009). Losing the null subject. A contrastive study of (Brazilian) Portuguese and (Medieval) French. In *Proceedings of the Workshop "Null-subjects, expletives, and locatives in Romance,"* 131–56.
- Kauhanen, Henri, & Walkden, George. (2017). Deriving the Constant Rate Effect. *Natural Language and Linguistic Theory*, 1–39.
- Koenenan, Olaf. (2006). Deriving the difference between full and partial pro-drop. In P. Ackema, P. Brandt, M. Schoorlemmer, & F. Weerman (Eds.), *Arguments and agreement*, 76–100. Oxford: Oxford University Press.
- Kratzer, Angelika. (2009). Making a pronoun: Fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry* 40:187–237.
- Kroch, Anthony. (1989). Reflexes of grammar in patterns of language change. *Language Variation and Change* 1:199–244.
- Manzini, Maria Rita & Savoia, Leonardo Maria. (2005). *I dialetti italiani e romanci: morfossintassi generativa*, volume 1. Edizioni dell'Orso.
- Marchello-Nizia, Christiane (1992). *Histoire de la langue française aux XIVe et XVe siècles*. Paris: Dunod.
- MCVF. (2010). Corpus MCVF annoté syntaxiquement, dirigé par France Martineau, avec Paul Hirschbühler, Anthony Kroch et Yves Charles Morin. URL http://www.voies.uottawa.ca/corpus_pg_en.html.
- Morin, Yves Charles. (2001). La troncation des radicaux verbaux en français depuis le Moyen Age. *Études diachroniques. Recherches linguistiques de Vincennes* 30:63–85.
- Nagy, Naomi & Heap, David. (1998). Franco-provençal null subjects and constraint interaction. *CLS* 34:151–66.
- Nagy, Naomi G., Aghdasi, Nina, Denis, Derek, & Motut, Alexandra. (2011). Null subjects in heritage languages: Contact effects in a cross-linguistic context. *University of Pennsylvania Working Papers in Linguistics* 17(2):134–44.
- Narendra, Kumpati S. & Thathachar, Mandayam A.L. (1989). *Learning Automata: An Introduction*. Englewood Cliffs, NJ: Prentice-Hall.
- Otheguy, Ricardo, Zentella, Ana Celia, & Livert, David. (2007). Language and dialect contact in Spanish in New York: Toward the formation of a speech community. *Language* 83:770–802.
- Paolillo, John C. (2011). Independence claims in linguistics. *Language Variation and Change* 23:257–74.
- Penn Supplement to MCVF. (2010). Penn Supplement to the MCVF Corpus by Anthony Kroch & Beatrice Santorini.
- Pintzuk, Susan. (1995). Variation and change in Old English clause structure. *Language Variation and Change* 7:229–60.
- Prévost, Sophie. (2018). Increase of Pronominal Subjects in Old French: Evidence for a Starting-point in Late Latin. In A. Carlier & C. Guillot (Eds.), *Latin tardif-français ancien: continuités et ruptures*, 169–98. Coll. Beihefte zur Zeitschrift für romanische Philologie. Berlin: De Gruyter.

- 1171 Ranson, Diana L. (2009). Variable subject expression in Old and Middle French prose texts: The role of
 1172 verbal ambiguity. *Romance Quarterly* 56:33–45.
- 1173 Rizzi, Luigi. (1986). Null objects in Italian and the Theory of *pro*. *Linguistic Inquiry* 17:501–57.
- 1174 Roberts, Ian. (1993). *Verbs and Diachronic Syntax: A Comparative History of English and French*.
 Dordrecht: Kluwer.
- 1175 ———. (2014). Taraldsen’s Generalization and Language Change: Two Ways to Lose Null Subjects. In
 1176 P. Svenonius (Ed.), *Functional Structure from Top to Toe: The Cartography of Syntactic Structures*,
 9, 115–48. Oxford: Oxford University Press.
- 1177 Santorini, Beatrice. (1993). The rate of phrase structure change in the history of Yiddish. *Language*
 1178 *Variation and Change* 5:257–83.
- 1179 Sauerland, Uli. (2008). On the semantic markedness of phi-features. In D. Harbour, D. Adger, & S. Bejar
 1180 (Eds.), *Phi Theory: Phi-features across Modules and Interfaces*, 527–82. Oxford: Oxford University
 Press.
- 1181 Schøsler, Lene. (2002). La variation linguistique: le cas de l’expression du sujet. In R. Sampson &
 1182 W. Ayres-Bennett (Eds.), *Interpreting the History of French, A Festschrift for Peter Rickard on*
the occasion of his eightieth birthday, 187–208. New York: Rodopi.
- 1183 Shlonsky, Ur. (2009). Hebrew as a partial null-subject language. *Studia linguistica* 63:133–157.
- 1184 Simonenko, Alexandra, Crabbé, Benoit, & Prévost, Sophie. (2015). Morphological triggers of syntactic
 1185 changes: Treebank-based information theoretic approach. In M. Dickinson, E. Hinrichs, A. Patejuk, &
 1186 A. Przepiórkowski (Eds.), *Proceedings of the Fourteenth International Workshop on Treebanks and*
Linguistic Theories (TLT14), 194–205. Warsaw: Institute of Computer Science, Polish Academy of
 Sciences.
- 1187 Speas, Margaret. (1995). *Economy, agreement and the representation of null arguments*. Ms. University
 1188 of Massachusetts, Amherst.
- 1189 Sprouse, Rex, & Vance, Barbara. (1999). An explanation for the decline of null pronouns in certain
 1190 Germanic and Romance dialects. In M. deGraff (Ed.), *Language creation and language change:*
creolization, diachrony and development, 257–84. Cambridge, Mass.: MIT Press.
- 1191 Taraldsen, Tarald. (1980). *On the NIC, vacuous application and the that-trace filter*. Indiana University
 1192 Linguistics Club.
- 1193 Tvica, Seid. (2017). On minimal person and number features. Unpublished manuscript. [https://www.
 1194 academia.edu/23278325/On_the_universality_of_person_and_number](https://www.academia.edu/23278325/On_the_universality_of_person_and_number).
- 1195 van Reenen, Pieter & Schøsler, Lene. (1987). Le problème de la prolifération des explications. *Vrije*
 1196 *Universiteit working papers in linguistics* 27.
- 1197 Vance, Barbara. (1997). *Syntactic change in medieval French*. Studies in Natural Language and
 1198 Linguistic Theory. Dordrecht/Boston/London: Kluwer.
- 1199 Vennemann, Theo. (1975). An explanation of drift. In Ch. N. Li (Ed.), *Word Order and Word Order*
 1200 *Change*, 269–305. Austin: University of Texas Press.
- 1201 Yang, Charles. (2000). Internal and external forces in language change. *Language Variation and*
 1202 *Change* 12:231–50.
- 1203 ———. (2002). *Knowledge and learning in natural language*. Oxford: Oxford University Press.
- 1204 ———. (2010). Three factors in language variation. *Lingua* 120:1160–77.
- 1205 Zimmermann, Michael. (2009). On the evolution of expletive subject pronouns in Old French. In E.-M.
 1206 Remberger & G.A. Kaiser (Eds.), *Proceedings of the workshop “null-subjects, expletives, and*
 1207 *locatives in romance.” Arbeitspapier 123*, 63–92. Fachbereich Sprachwissenschaft, Universität
 1208 Konstanz.
- 1209 Zimmermann, Michael. (2014). *Expletive and referential subject pronouns in Medieval French*, volume
 1210 556. Berlin, Boston: Walter de Gruyter GmbH.
- 1211
- 1212
- 1213
- 1214
- 1215