

Investigating the Evolution of Agreement Systems Using an Artificial Language Learning Paradigm

Jennifer Culbertson and Géraldine Legendre
Johns Hopkins University

1. Introduction

The evolution of new agreement systems from grammaticalized pronominal elements is one of the most well-documented processes of linguistic change (Givón 1976; Hopper and Traugott 1993; van Gelderen 2004; Fuss 2005). This pathway can be described by the grammaticalization cline in (1). This paper targets the final stages of this cline, which we hypothesize to be constrained by an implicational hierarchy. A simplified version of this hierarchy, alternately called the Topicality, Definiteness, or Accessibility Hierarchy (Givón 1976; Ariel 2000; Siewierska 2004), is shown in (2).

(1) *Grammaticalization cline (Hopper and Traugott 1993)*
independent pronoun → clitic pronoun → agreement affix

(2) *Definiteness Hierarchy*
Pronoun > Definite DP > Indefinite DP

We focus on two critical observations about the evolution of new agreement markers: (i) they start out as optional, and gradually become obligatory, and (ii) they are used first with DP types on the left of the Definiteness Hierarchy and are generalized rightward (step-by-step). We argue that the process of making agreement obligatory with a given DP type—*regularization*—and the process of extending agreement to new DP types—*over-generalization*—are driven by the learner, and constrained by the Definiteness Hierarchy.

To show this, we first provide evidence from the ongoing evolution of new agreement in French. Then we report the results of an artificial language learning experiment modeled after French to directly test our hypotheses. We show that learners in the experiment regularize and over-generalize variable

systems of agreement, but that their behavior is crucially dependent on whether the language is predicted possible by the Definiteness Hierarchy.

2. A case of ongoing change: the French agreement system

French subject clitics fall somewhere between clitic pronouns and agreement affixes—some authors analyze them as syntactically independent arguments, cliticized phonologically (e.g. Kayne 1975; Rizzi 1986; De Cat 2007), while others treat them as affixal agreement markers (e.g. Auger 1994; Culbertson 2010a). Culbertson (2010a) argues that while in the more conservative Standard French register, they may still function as argumental subjects, in the Colloquial French register they are clearly agreement affixes. That these elements are affixes rather than phonological clitics is based on morphophonological and distributional evidence. Evidence that they function as agreement markers comes from their co-occurrence with pronominal and lexical DP subjects in so-called *subject doubling* constructions. An example is shown in (3).

- | | |
|------------------------------|---|
| (3) Lui/Jean il parle. | (4) *Personne il parle. |
| him/Jean scl talks | nobody scl talks |
| 'Him, Jean (he) is talking.' | 'Nobody is talking.' (intended meaning) |

Although French subject doubling has been traditionally analyzed as topic-left-dislocation (the clitic acting as a resumptive pronoun), Culbertson (2010a) shows, on the basis of experimental evidence, that these constructions are not prosodically marked as dislocation, and can occur even when the lexical subject is in-situ. Further, Culbertson (2010a) shows that subject doubling is pervasive in several adult- and child-directed speech corpora of Colloquial French, suggesting it is not a pragmatically marked construction.

However, even if Colloquial French subject clitics have transitioned from cliticized subjects to affixal agreement markers, this new system of agreement still has room to evolve. In particular, while (3)—doubling of a definite DP—is not only possible, but frequently attested, (4)—doubling of an *indefinite DP*—is impossible. Colloquial French in fact falls in the middle of the Definiteness Hierarchy in (2); subject doubling is obligatory with pronominal DP subjects, optional with definite lexical DPs, and impossible with indefinite lexical DPs.

To summarize: in the linguistically conservative Standard French register, subject clitics are relatively loosely bound to the verb, and act like argumental subject pronouns. In the Colloquial French register, on the other hand, these elements behave like tightly bound affixal agreement markers, but their appearance depends on the definiteness of the subject DP. It is easy to imagine a future change in which these new agreement markers occur with indefinite subject DPs as well (in fact Picard, a language closely related to French, has undergone this further change, see Auger 2003). Why is it the case that as a

weak pronominal element like the French subject clitic becomes increasingly bound to its host verb, it also tends to be used with greater regularity? The explanation we are suggesting is that language learners play an important role in driving this change, altering the target grammar by using the element more often (regularizing), and with fewer contextual restrictions (over-generalizing).

Evidence that child learners have a tendency to regularize inconsistent variation in their language can be found in work on creolization, (Sankoff and Laberge 1980; Sandler et al 2005), and acquisition of multiple registers (e.g. Smith et al 2007). In the case of French, the fact that subject doubling constructions as in (3) are used more frequently by younger speakers suggests that new generations of speakers are in fact pushing the language along the diachronic path. Figure 1 shows a best-fit linear regression line of doubling rate by age for speakers in three corpora of adult-directed speech.¹

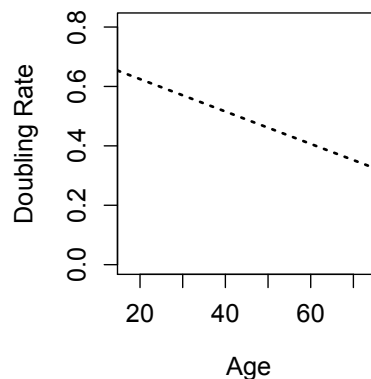


Figure 1. Doubling rate by age in the Coveney, PFC, and CID corpora ($r^2=0.46$)

Data from French also provide evidence that learners may be responsible for over-generalizing emerging agreement systems. French children's spontaneous speech contains intriguing evidence of non-adult-like use of subject doubling which follows precisely the pattern predicted by the Definiteness Hierarchy. Specifically, they also produce subject doubling with clearly indefinite DP subjects—impossible in the adult grammar. Several examples from the Lyon Corpus (Demuth and Tremblay 2008) are shown below in (5).

- (5) a. Un couteau il est là. (Ana 2;9) b. Un cube il est tombé. (Tim 2;1)
 a knife scl is there a block scl is fell
 '*A knife is there.*' '*A block fell.*'

Do learners drive the evolution of new agreement systems? Are they responsible for shifting languages along the implicational hierarchy of Definiteness? Above we have provided some suggestive evidence that French

learners may in fact advance the language they acquire along this diachronic pathway. However, the strength of the conclusions we can draw from this evidence is limited.

Below we report the results from an artificial language learning experiment, designed to uncover clear evidence of the potential role of the learner in this type of apparently constrained morphosyntactic change. The experiment investigates a formal bias on the part of the learner to regularize variation present in the input—the *regularization bias*—as well as a *substantive bias* favoring grammars in line with the Definiteness hierarchy. The experiment reveals that learners reliably regularize the variation present *only* when the input follows the hierarchy, when it does not, they shift towards a grammar which better satisfies the substantive bias, rather than regularize the input pattern.

3. Experimental evidence for the role of the learner

The experiment reported here uses an artificial language learning paradigm—similar to that developed in Hudson Kam & Newport (2009)—which takes advantage of the hypothesis that learners will regularize inconsistent variation found in the input. This paradigm is a logical choice to study biases on the part of the learner relevant to language change generally, and the evolution of new agreement systems specifically, since both involve the introduction of variation followed by regularization of a new pattern. Following the basic methodology used in Hudson Kam and Newport (2009), the artificial languages feature a *grammar mixture* that is skewed toward a general pattern, used the majority of the time, but accompanied by some noise. We use this paradigm to explicitly test learners’ willingness to regularize linguistic patterns which are predicted to differ based on hypothesized substantive biases.

3.1. Design, methodology and predictions

This experiment exposes learners to a language with a variably present element whose phi-features (here gender and number) match those of the grammatical subject. We label this agreeing element a *clitic* as a descriptive convenience, since its properties are designed to parallel those of the French subject clitic. Following the hypothesis that new agreement systems evolve (in part) through regularization by language learners, the experiment examines the conditions under which learners exhibit a tendency to use a variable element—in this case the clitic—more often than it is present in the input.²

As mentioned above, we are also interested in how learning is influenced by the Definiteness Hierarchy in (2). The hierarchy predicts that a language allowing agreement with definite DPs but not indefinite DPs is possible (Colloquial French is such a language). By contrast, a language allowing agreement with indefinite DPs but not definite DPs is predicted to be impossible

(and is cross-linguistically unattested). By hypothesizing an interaction between the bias to regularize and the substantive bias—the Definiteness Hierarchy—we predict that learners will regularize *only* a language which accords with this hierarchy.

3.2. Lexicon and sentence types

Sentences in the miniature artificial language describe actors, either definite or indefinite, performing some action. The lexicon consists of 4 intransitive verbs (*voitch* ‘kick’, *ramza* ‘wave’, *cherg* ‘sit’, *geeja* ‘put hands on face’), 6 definite nouns (*zadgi* ‘frog’, *blifu* ‘lion’, *maugi* ‘monkey’, *nerki* ‘lemur’, *trefu* ‘opossum’, *flarmu* ‘ant’)—half feminine (ending in *-u*), and half masculine (ending in *-i*)—and 2 indefinite nouns (fixed forms, *griftor* ‘everybody’, *slergor* ‘some’). Table 1 shows the paradigm of the clitic, which never bears stress or is separated from the verb, thus resembling a typical clitic or affix. Table 1 also provides the paradigm for two definite DPs. Each is marked for gender, and number (the plural marker is *-la*). The vowel alternation indicating the gender of the clitic is identical to the alternation in the set of masculine and feminine definite DPs. Verbs do not have any additional morphological marking.³

Clitic	Features	DP	Features
si	MASC, SG	nerk-i	MASC, SG
su	FEM, SG	tref-u	FEM, SG
di	MASC, PL	nerk-i-la	MASC, PL
du	FEM, PL	tref-u-la	FEM, PL

Table 1. Artificial language lexicon; clitic and definite DP paradigms.

The grammar generates three sentence types, shown in (7).

(7) *Artificial language sentence types (examples)*

- (i) S → DP Verb (nerki geeja. ‘The lemur waves’)
- (ii) S → DP Clitic Verb (nerki si geeja. ‘The lemur clitic-waves.’)
- (iii) S → Clitic Verb (si geeja. ‘(He) clitic-waves.’)

The DP-types (definite and indefinite) which appear in sentences of types (i) and (ii) vary across conditions (the precise manipulation is discussed below). Sentences of type (iii) are used only when the actor depicted is definite regardless of condition (just as would be the case in French).⁴

The artificial languages created with this lexicon and set of sentence types are designed to be asymmetrical with respect to the DP-types used; as in natural languages, most sentences have definite as opposed to indefinite subjects.⁵ While this feature makes the artificial languages more realistic, it also introduces the possibility that definite DPs could be learned more successfully than

indefinite DPs. To compensate for this asymmetry, each of the 6 definite and 2 indefinite referents were used as the actor(s) in a sentence an equal number of times across training. Since sentences of type (iii), without an overt DP subject, are always definite, this means each individual definite DP will actually be heard *overtly* fewer times than each individual indefinite DP.

3.3. Conditions

Participants in the experiment were randomly assigned to one of two conditions, illustrated in Table 2. In the first, overt definite DP subjects are optional agreement triggers; when an overt definite DP subject is present, it triggers a clitic 75% of the time. Indefinite DP subjects never appear with the clitic. In the second condition, the opposite is true—indefinite DP subjects are optional agreement triggers and when they are present they trigger the clitic 75% of the time. Definite DP subjects never appear with the clitic. The conditions differ critically according to the Definiteness Hierarchy—a language in which definite but not indefinite DPs trigger agreement is perfectly plausible. However, the opposite language, namely one in which clitics are triggered by indefinite but not definite DPs is *impossible*. The first condition is thus labeled the "natural" condition, while the second condition is labeled "unnatural".

	Natural Condition		Unnatural Condition	
	Definite	Indefinite	Definite	Indefinite
(i) S → DP Verb	75%	100%	100%	25%
(ii) S → DP Clitic Verb	25%	0%	0%	75%

Table 2. Conditions; note that this does not include type (iii) sentence, in which no overt DP is present, which are constant across conditions.

3.4. Hypotheses and predictions

The hypotheses this experiment is designed to test, shown in (8), are centered around the idea that learners will generally tend to regularize variation, but are constrained by substantive biases. During the learning process, both biases may cause learners to acquire a grammar which differs systematically from the input. In this section we will outline how these hypotheses can be used to derive predictions about how learners will behave in this experiment.

(8) *General hypotheses*

1. Learners have a *regularization bias*, favoring less variable grammars
2. Learners have a *substantive bias*, favoring grammars which are allowed by the Definiteness Hierarchy.
3. The substantive bias outweighs the regularization bias.
4. Over-generalization is more costly than regularization.

By positing an interaction between substantive bias and the regularization bias, it is possible to generate predictions about learners' behavior given some variable rule in the input grammar. First, learners should regularize the use of a variably present agreeing clitic when the input respects the Definiteness Hierarchy. Learners in the natural condition are therefore predicted to regularize the use of the clitic element with definite DPs. They may also over-generalize, but we assert this is more costly than regularization.⁶ When the input does *not* respect the hierarchy—as in the unnatural condition—the substantive and regularization biases are in conflict, and following Hypothesis 3, the prediction is that learners will *not* regularize.

How might learners in the unnatural condition shift the grammar to bring it *more* in line with the Definiteness Hierarchy? They could move towards a grammar in which *no* DPs trigger the clitic, or towards a grammar in which *all* DPs trigger it. We predict that movement toward the latter is less probable than movement toward the former, since the latter involves over-generalization. (Movement towards a grammar in which definite but not indefinite DPs trigger agreement can be ruled out; this type of shift both fails to regularize *and* involves over-generalization.) These predictions are laid out in (9).

(9) *Predictions*

1. Learners in the natural condition will regularize, may over-generalize.
2. Learners will not regularize in the unnatural condition.
3. Learners in the unnatural condition will shift toward a language which never uses clitics, or uses them will all DP types.

3.5. General procedure

As mentioned above, this experiment is roughly modeled after the method developed in Hudson Kam and Newport (2009). Participants were trained and testing on a miniature artificial language during a single hour-long session using a videogame interface. They were seated in a private testing room in front of a computer display, and wore headphones through which the experiment audio was played. Participants were introduced to a native speaker informant whose utterances they learned from.⁷

Each session involved two phases of exposure training (200 trials) followed by comprehension and production testing (80 trials each). Throughout, the informant uttered phrases stochastically according to the probabilities specified by the condition the participant was assigned to. Grammars were probabilistic in the sense that given a scene and the vocabulary items which could be used to describe it, the particular structural realization was determined only by the relative probabilities assigned to the rule expansions in the grammar.

Here we present results from the production testing phase (results from other phases are reported in Culbertson 2010b). In this phase, participants were shown

a scene, and were instructed to provide a description of that scene in the language. To ensure that in providing their description participants would have to choose between the critical sentence types (i) and (ii), the informant provided a DP, and participants were told to use it in the sentence.

3.6. Participants

Participants in the experiment were 32 native English speaking undergraduates at Johns Hopkins University (16 in each condition; 7 males, 18-25 years old). One additional participant failed a preliminary vocabulary learning task and was excluded. They received either course credit or \$15 for their participation.

3.7. Experiment results and analysis

Figure 2 shows the proportion of participants' utterances which used the clitic with definite and indefinite DP subjects, by condition. As the figure suggests and a one-sample Sign-Test confirms, learners in the natural condition *regularized*—they used clitics with agreement triggering definite DPs more than 75% of the time ($p=0.04$). By contrast, learners in the unnatural condition did not regularize ($p=0.91$), but rather produced clitics with agreement triggering indefinite DPs less than 60% of the time. For DP types which did not trigger clitics in the input—indefinite DPs for the natural condition, and definite DPs for the unnatural condition—95% binomial confidence intervals were constructed for each condition using the mean level of over-generalization.⁸ For both conditions, the 95% confidence intervals did not include 0% (for the natural condition, $5\% \leq \mu \leq 55\%$; for the unnatural condition, $3\% \leq \mu \leq 62\%$), confirming that learners in both conditions over-generalized.

To compare performance across conditions, the data were subjected to mixed-effects logistic regression. A model with clitic use as the dependent variable, participant as a random effect, and condition and DP-type as fixed effects, revealed a significant effect of condition ($\beta = -1.27$, $z = -5.91$, $p < 0.0001$) and DP-type ($\beta = -0.92$, $z = -5.97$, $p < 0.0001$), as well as an interaction ($\beta = 2.7$, $z = 17.62$, $p < 0.0001$). This indicates that, (i) overall, learners in the unnatural condition used significantly fewer clitics compared to the natural condition, (ii) overall, learners used clitics significantly less often with indefinites than with definites, and (iii) learners in the unnatural condition were significantly more likely to use clitics with indefinite DPs than learners in the natural condition.

To further investigate these differences, individual planned contrasts were also tested. To confirm that the natural and unnatural conditions differed from one another with respect to regularization of the clitic with agreement-triggering DPs, a mixed effects model with use of a clitic as the dependent variable, and condition as a fixed effect was fit to the relevant subset of the data. This model revealed a significant difference between the conditions with respect to use of

the clitic with agreement-triggering DPs; learners in the unnatural condition used fewer clitics with agreement-triggering DPs than learners in the natural condition did ($\beta=-0.713, z=-8.81, p<0.0001$). The level of over-generalization in both conditions was also compared, revealing no significant difference between conditions ($\beta=-0.112, z=1.34, p=0.18$).




Figure 2. Use of clitic by DP type by condition in the production test.

The results so far reported clearly confirm the predictions in (9). Learners in the natural condition—in which the input pattern conformed to the Definiteness Hierarchy—exhibited regularization. Consistent with our predictions learners also over-generalized clitics to non-agreement-triggering DPs.

As for the unnatural condition—in which the input pattern did not conform to the Definiteness Hierarchy—as predicted, learners did *not* regularize. As a group, learners in the unnatural condition, whose input mixture featured only indefinite DPs as agreement triggers, shifted towards a distribution of grammars which less strongly violates the substantive bias. Specifically, the discrepancy between the use of the clitic with the two DP types in the input was reduced, with learners using fewer clitics with indefinite DPs and more clitics with definite DPs (as compared to the input).

Because more than one strategy for shifting the input mixture was predicted possible, we now turn to individual learner outcomes. Figure 3 is a plot of the output probabilities for each DP-type for individual participants. The x-axis

shows the probability of using a clitic given definite DP subject, and the y-axis the probability of using a clitic given an indefinite DP subject. Each of the corners in the plot represents one of four logically possible deterministic grammars: G1, in which clitics are never used, G2, in which they are always used, G3, in which clitics are triggered by definite DPs only, and G4, the “unnatural” grammar in which clitics are triggered by indefinite DPs only.




Figure 3. Plot of individual learner outcomes. Asterisks are training points, circles are individual testing points.

The dashed circles in the G3 and G4 corners of the plot are where individual learners should fall if they are regularizing. Supporting the predictions, and confirming the conclusions drawn above based on mean performance, learners in the natural condition were very likely to shift toward the G3 corner—falling within the circle surrounding that corner. A number of learners not only regularized but over-generalized as well, shifting toward the G2 corner. However, no learners fell into the circle surrounding the G4 corner, rather learners in the unnatural condition overwhelmingly moved toward the G1 and G2 corners, as predicted.

4. Discussion

When weakened pronominal clitics become new agreement markers, they start out by appearing only variably with triggering DPs. Over time, they are

regularized and extended to new DP types. In the case of French, the use of subject doubling constructions involving lexical DP subjects and subject clitics, has increased over generations, and has spread from pronominal to definite DPs (as in (3) above). This construction, which once involved left-dislocation, has become an agreement dependency (Culbertson 2010). The role of the learner in change along this trajectory is important from the perspective of both diachrony and language acquisition. The biases learners exhibit are crucial to understanding how language change is actuated and how changes spread. At the same time, evidence for both substantive and formal learning biases (like the regularization bias) reveal both structural and formal preferences on the part of the learner that inform our theories of the knowledge the learner brings to the acquisition task.

Although we have some evidence supporting the role of French learners in this particular case, here we have used an artificial language learning experiment to provide clear behavioral evidence for our hypotheses. Learners in the experiment reported here were exposed to a language in which a clitic-like element, whose features agreed with those of the sentential subject, appeared only optionally, and co-occurred with only some types of DP subjects. We were able to show that when the input language respected the Definiteness Hierarchy, learners tended to regularize the variation present in the input. Similar to what is found with child learners of French, learners also over-generalized the agreeing clitic to new types of DP subjects, namely indefinites. This result bolsters support for the claim that regularization and extension of new agreement systems are driven in part by biases on the part of the learner to increase the regularity in grammatical systems.

Critically, however, when learners were exposed to a typologically *unnatural* language—a language predicted impossible by the Definiteness Hierarchy hypothesized to constraint agreement systems—they behaved quite differently. In particular, learners did *not* regularize, but shifted the language, bringing it in line with the Definiteness Hierarchy—exactly as predicted if the hierarchy acts as a substantive bias constraining acquisition. Although change in natural languages undoubtedly involves many factors, the differences in behavior found across conditions in the experiment suggests that biases on the part of the learner may play a critical role. Further, these differences strongly suggest that grammar learning involves more than tracking probabilities; in this case, learning is crucially influenced by substantive constraints on agreement patterns.

The artificial languages learners were exposed to here incorporate an asymmetry in the type (although not token) frequency of definite and indefinite DP subjects—maintaining a realistic feature of natural languages including French. This leaves open the possibility that the differences found in the use of the clitic across conditions was due in some part to this asymmetry. If this effect can be ruled out (see Culbertson 2010b for some evidence that it can), then the results of this experiment provide strong evidence that language-specific biases

need to be incorporated into theories of language acquisition to account for typological constraints on linguistic patterns (contra Evans and Levinson 2009; Goldberg 2006, and others). This will be addressed in detail in future work.

Notes

- ¹ The linear regression model has clitic doubling rate as the dependent variable, and age and corpus as fixed effects. Age is a significant factor ($p < 0.001$) in the model.
- ² By investigating the extent to which participants will regularize this kind of variation, this experiment is similar to Hudson Kam and Newport (2009). However, here we are interested in testing regularization of a variable *agreement dependency*—that is, regularization of a morphological element whose form varies according to the features of the sentential subject).
- ³ The clitic provides the only morphological dependency between the subject and verb, similar to Colloquial French where suffixal morphology for about 90% of verbs is not phonologically distinct.
- ⁴ These sentences are potentially important—learners might be more likely to form a representation of the clitic and the verb as a single unit if they hear a large number of such sentences, and therefore at the point of lexical retrieval might be more likely to retrieve the clitic+verb unit rather than the verb alone. This cohesive representation could encourage learners to use of the clitic with DP subjects more, (this is itself a hypothesis, however it is assumed here rather than explicitly tested).
- ⁵ Indefinite subjects are rare cross-linguistically, and in fact some languages do not allow them at all (for general discussion see Comrie 1989).
- ⁶ While regularization and over-generalization both alter the input grammar, the latter takes a DP-type whose probability of co-occurring with a DP is *zero* and increases it.
- ⁷ The informant's speech was synthetically generated using Apple's text-to-speech software (OS 10.5, speaker "Alex").
- ⁸ Since clitics are never used with these DP types in the input, a Sign-Test cannot be used to test for over-generalization. Confidence intervals can be used instead to estimate a range of values which are likely to include some population parameter, like the mean, with some level of confidence (here 95%). If the range for the population parameter does not include that value, we can be 95% confident that this is not the true parameter value. The data in this case are proportions, so we use binomial confidence intervals constructed using the Wilson score interval (Agresti and Coull 1998).

References

- Agresti, A. and B. Coull. 1998. "Approximate is better than 'exact' for interval estimation of binomial proportions." *American Statistician*, 52:119–126.
- Ariel, Mira 2000. "The development of person agreement markers: from pronouns to higher accessibility markers." *Usage-Based Models of Language*, ed. Michael Barlow and Susan Kemmer, 197–260. Stanford: CSLI Publications.
- Auger, Julie 1994. *Pronominal Clitics in Québec Colloquial French : A Morphological Analysis*. Ph.D. thesis, University of Pennsylvania.
- Auger, Julie. 2003. "Les pronoms clitiques en picard: une analyse au confluent de la phonologie, de la morphologie et de la syntaxe." *Journal of French Language Studies*, 13:1–22.
- Comrie, Bernard. 1989. *Language universals and linguistic typology*. Chicago: University of Chicago Press.

- Culbertson, Jennifer. 2010a. "Convergent evidence for categorical change in French: from subject clitic to agreement marker." *Language*, 86(1): 85-132.
- Culbertson, Jennifer. 2010b. Learning biases, regularization, and the emergence of typological universals in syntax. Ph.D. thesis, Johns Hopkins University.
- De Cat, Cécile 2007. *French Dislocation. Interpretation, Syntax, Acquisition*. New York: Oxford University Press.
- Demuth, Katherine and Annie Tremblay. 2008. "Prosodically-conditioned variability in children's production of French determiners." *Journal of Child Language*, 35(1):99-127.
- Evans, Nicholas and Stephen Levinson. 2009. "The myth of language universals: Language diversity and its importance for cognitive science." *Behavioral and Brain Sciences*, 32(05):429-448.
- Givón, Talmy 1976. "Topic, Pronoun, and Grammatical Agreement." *Subject and Topic*, ed. Charles Li, 149-188. New York: Academic Press.
- Goldberg, Adele. 2006. *Constructions at Work: The nature of generalization in language*. New York: Oxford University Press.
- Hopper, Paul and Elizabeth Traugott. 1993. *Grammaticalization*. New York: Cambridge University Press.
- Hudson Kam, Karla and Elissa Newport. 2009. "Getting it right by getting it wrong: When learners change languages." *Cognitive Psychology*, 59(1):30-66.
- Kayne, Richard. 1975. *French Syntax: The Transformational Cycle*. Cambridge: MIT Press.
- Rizzi, Luigi. 1986. "On the status of subject clitics in Romance." *Studies in Romance Linguistics*, eds. Osvaldo Jaeggli and Carmen Silva-Corvalan, 391-419. Dordrecht: Foris
- Sandler, Wendy, Carol Padden, and Mark Aronoff. 2005. "The emergence of grammar: Systematic structure in a new language." *PNAS*, 102(7):2661-2665.
- Sankoff, Gillian & Suzanne Laberge. 1980. "The acquisition of native speakers by a language." *The Social Life of Languages*, ed. Gillian Sankoff, 195-209. Philadelphia: University of Pennsylvania Press.
- Siewierska, Anne. 2004. *Person*. New York: Cambridge University Press.
- Smith, Jennifer, Mercedes Durham, and Liane Fortune. "‘‘mam, my trousers is fa’in doon!’’: Community, caregiver, and child in the acquisition of variation in a Scottish dialect." *Language Variation and Change*, 19:63-99.