

Syntactic production is not independent of inhibitory control: Evidence from agreement attraction errors

Nazbanou Nozari (nozari@jhu.edu)

Department of Neurology, Department of Cognitive Science, Johns Hopkins University
1629 Thames St., Suite 350, Baltimore, MD 21231, USA

Akira Omaki (omaki@uw.edu)

Department of Linguistics, University of Washington
3940 Benton Ln NE, Box 352425, Seattle, WA 98105, USA

Abstract

Native adult speakers of a language can produce grammatical sentences fluently, effortlessly, and with relatively few errors. These characteristics make the highly-practiced task of speaking a viable candidate for an automatic process, i.e., one independent of cognitive control. However, recent studies have suggested that some aspects of production, such as lexical retrieval and tailoring speech to an addressee, may depend on the speaker's inhibitory control abilities. Less clear is the dependence of syntactic operations on inhibitory control processes. Using both a direct manipulation of inhibitory control demands and an analysis of individual differences, we show that one of the most common syntactic operations, producing the correct subject-verb agreement, requires inhibitory control when a singular subject noun competes with a plural local noun as in "The *snake* next to the purple *elephants* is green." This finding calls for the integration of inhibitory control mechanisms into models of agreement production, and more generally into theories of syntactic production.

Keywords: syntactic production; subject-verb agreement; agreement attraction; inhibitory control; individual differences

Introduction

The last few decades of sentence production research have shed much light on the nature of the grammatical encoding mechanisms that convert abstract thoughts into multiple levels of linguistic representations, such as words, syntactic structures, and phonetic output (e.g., Levelt, 1989). In comparison, little attention has been paid to whether and how non-linguistic cognitive operations, such as inhibitory control, support such mechanisms. Perhaps the main reason for this neglect is that language production, at least in native adult speakers, is a well-practiced, highly efficient, and reasonably error-free process. In addition, after the conceptual message is constructed, relatively little conscious effort goes into the subsequent processes of lexical retrieval, phonological encoding, and building a syntactic frame. All of these attributes make language production a viable candidate for an "automatic" process, i.e., a process that operates independently of cognitive control (e.g., Shiffrin & Schneider, 1977).

Recently, however, more evidence has come to light pointing to the relevance of cognitive control processes, especially inhibitory control, to the various aspects of language production such as lexical retrieval (e.g., Shao,

Meyer, & Roelofs, 2013) or accommodating listener's perspective (Trude, & Nozari, 2017). One area of sentence production in which the potential role of inhibitory control has remained understudied is syntactic operations. As a consequence, syntactic theories have remained, for the most part, disconnected from cognitive theories. This study proposes a first step to reconcile the two by investigating the role of inhibitory control in one of the most basic syntactic processes in English and many other languages, namely the process of subject-verb agreement, where speakers must select the verb form that agrees in number with its subject. For native adult speakers, agreement production is effortless and mostly error-free. However, it is not uncommon for sentences to contain more than one noun, in which case the subject noun (N1) could compete with the local noun (N2) for determining verb agreement. For example, "The *snake* next to the *elephants*..." could elicit the verb "are" because N2 is plural, even though N1 is singular. This phenomenon is called "agreement attraction" (Bock and Miller, 1991). Figure 1 shows the syntactic structure subject to agreement attraction.




Figure 1: Syntactic structure subject to agreement attraction.
sg = singular; pl = plural.

Broadly speaking, two classes of models explain the agreement attraction phenomenon for the configuration shown in Figure 1. The first class (e.g., *Feature Percolation, Marking and Morphing*) attributes agreement attraction to faulty representation of number information on the subject NP, which is hypothesized to result from additional processes such as feature percolation (e.g., Franck,

Vigliocco, & Nicol, 2002) or spreading activation (e.g. Eberhard, Cutting, & Bock, 2005). The second class (e.g., *Memory Interference*; Badecker & Kuminiak 2007) attributes attraction effects to constraints on content-addressable working memory mechanisms that result in occasional retrieval of N2 for subject-verb agreement. None of these accounts currently propose a role for inhibitory control.

The current study

While none of the models of agreement attraction mentioned above explicitly propose a role for inhibitory control in the agreement attraction process, all of them entail processes that lend themselves well to an influence of inhibitory control. Table 1 provides a summary. In light of this, it seems reasonable, and we argue essential, to test whether inhibitory control indeed plays a role in preventing agreement attraction errors from arising. If so, then a complete theory of agreement production must include a mechanism like those proposed in Table 1. More importantly, impoverished inhibitory control should be taken into account when studying production problems such as high rates of agreement attraction errors in children (e.g., Franck et al., 2002), and possibly other populations with impaired inhibitory control such as individuals with aphasia. The absence of any evidence that inhibitory control is critical for preventing agreement attraction errors, in turn, lends credibility to claims of automaticity in language production.

We investigated the potential role of inhibitory control in agreement attraction in two ways: (a) by direct manipulation of the need for inhibitory control and examining its consequences on the rate of agreement attraction errors, and (b) by an analysis of individual differences to test whether the variability in individuals' scores on inhibitory control tasks predicts the variability in their agreement attraction errors in a production task. Most studies of agreement attraction use a preamble paradigm, in which participants hear an incomplete sentence (preamble) such as "The key next to the cabinets..." and must repeat and complete the sentence with an ending of their choice. While useful for eliciting agreement errors, this paradigm does not really capture the natural processes involved in everyday sentence production. Speakers rarely have to memorize and repeat

fragments while assembling an ending for a sentence they just heard out of context, without access to the full message. This is particularly relevant when comparisons are made between child and adult patterns of speech errors, as children are much less likely to be able to adapt to the unusual demands of the preamble paradigm.

We thus chose a referential communication paradigm in which task goals naturally translate into production of sentences with a structure suitable for studying agreement attraction. Participants described colored animals to a confederate so she could color gray animals on her own screen (see Figure 1). There were always more than one animal of the same type on the confederate's screen, so the participants learned that in order to unambiguously describe the target animal's color (a snake that is green) to the confederate they had to use the cue animals and produce a sentence such as "The snake *next to the purple elephant* is green". These targets and cues were arranged to create the four trial types summarized in Table 2.

To manipulate inhibitory control in this task, we took advantage of an effect reported by Gleitman et al. (2007). In that study, participants had to describe visual scenes with two objects, one of which was visually cued on each trial. The authors found that participants had a reliable tendency to produce the cued object first. For example, if the dog was cued first, they produced "The dog chased the cat". When the cat was cued, they produced "The cat was chased by the dog". The robust preference for thematic role assignment such that the cued object is produced first implies that overriding such a preference requires inhibitory control.

Inspired by Gleitman et al. (2007), we manipulated the need for applying inhibitory control by including two conditions: in the Target-flash (control) condition, participants were instructed to describe the color of the animal that was flashing on the screen. The sentence thus started with the visually-cued animal and no inhibition was required. In the Cue-flash (experimental) condition, participants were told that the flashing animal was the cue object, and that they had to describe the color of the animal next to it. Thus in this condition, speakers were forced to inhibit the preferred thematic role assignment which put the visually-cued animal first, in favor of one which allowed them to produce the non-cued animal first. If inhibitory control plays a role in preventing agreement attraction errors, then the Cue-flash condition which consumes some

Table 1: Existing models of agreement attraction, and a potential role for inhibitory control in each.

Theory	Attraction Mechanism	Potential effect of inhibitory control
Feature Percolation	The number feature of the local noun (N2) percolates to the subject NP (N1), and controls the verb form selection	Suppression of feature percolation
Marking and Morphing	Verb number is a function of notional number of N1, as well as values of the plural morphemes in N1 and the N2 each multiplied by the weight between the root and the morpheme	a) Decreasing the value of N2's morpheme b) Decreasing the weight of the link between morpheme and the root
Memory Interference	Agreement relies on retrieving and checking the number feature of what should determine agreement. Since N2 has partial cue-overlap with N1, it competes with it for selection.	Biasing the competition towards N1 and away from N2 (i.e., N2 suppression)

of the inhibitory control resources for the correct thematic role assignment should lead to more agreement attraction errors. On the other hand, if agreement production is automatic, the difficulty associated with inhibiting the assignment of subject role to the cued animal in the Cue-flash condition should have no effect on agreement attraction errors.

Table 2: Examples of the four trial types. S = Singular, P = Plural. The first letter denotes the plurality of N1; the second letter, the plurality of N2.

Trial type	Example
SS	The snake next to the purple elephant is green.
SP	The snake next to the purple elephants is green.
PP	The snakes next to the purple elephants are green.
PS	The snakes next to the purple elephant are green.

While a higher rate of agreement attraction errors in the Cue-flash compared to the Target-flash condition would imply the dependence of agreement production on cognitive control, it would be harder to conclude with certainty that it is specifically inhibitory control, and not a general attentional process to overcome task difficulty, that is required for agreement production. Therefore, we also report an analysis of individual differences in which we use a variant of structural equation modeling to extract the common variance from a battery of inhibitory control tests. We then examine whether the estimated inhibitory control ability of individuals predicts the variability in the rate of agreement attraction errors they produce.

Methods

Participants

Fifty-four native English speakers ($M_{age} = 22.80$, $SD = 4.03$ yrs.) completed two sessions in exchange for payment.

Materials and Procedures

Referential Communication Task All participants completed the referential communication task in session 1. Stimuli consisted of pictures of four animals (snake, elephant, lion, and bird), each appearing in two of four colors (purple, red, green, and brown; Figure 1). Animals were arranged in groups, containing a “target” (animal(s) whose color was to be described) and a “cue” (animal(s) used to help disambiguate the target animal). Both target and cue could be either one or two animals, creating four target-cue combinations: SS, SP, PP, and PS, where S = singular, P = plural (see Table 2). Participants completed two conditions, Target-flash and Cue-flash, each containing 128 trials (total = 256), in counterbalanced order. Each condition was divided into 4 blocks of 32 trials with breaks in between blocks. The 32 trials were presented in 8 slides. Figure 1 shows an example of a slide (left: participant view; right: confederate view). On each slide, six target-cue pairs

were presented, only four of which were to be described. At the beginning of each event, one of the animals (or animal pairs) flashed twice. In the Target-flash condition, this was the animal whose color was to be described to the confederate. At the beginning of the experiment, the confederate’s screen was shown to the participants, and it was explained that they should help the confederate color the gray animals. Participants learned that they could not just say “The snake is green”, because confederate’s screen contained two gray snakes. They thus had to use the cue animals to disambiguate their target by saying “The snake next to the purple elephants is green.” Once participants understood this logic, they naturally used the same sentence structure to describe other animals without having to memorize the structure. They completed 2 practice slides (8 trials), followed by 128 experimental trials, containing equal numbers (32) of SS, SP, PP, and PS trials. The order was arranged such that each trial type followed each of the other trial types with equal frequency.




Figure 1: Example slide from the referential communication task. Left: participant view. Right: confederate view. Yellow circle = flashing animal.

The Cue-flash condition was identical to the Target-Flash condition, except that participants were instructed that the flashing animal was the cue and the animal whose color was to be described was the one next to the flashing animal. For example, in Figure 1, instead of the snake, the elephants would flash, but the sentence would still be “The snake next to the purple elephants is green.” Participants completed 8 practice trials, followed by 128 experimental trials arranged according to the same criteria described for the Target-flash condition. Materials were displayed at the center of a 15-by-12 inch Dell monitor approximately 25 inches in front of the participants using the E-Prime 2.0 software (Psychology Software Tools, Pittsburg, PA). Participants’ speech was digitally recorded for offline transcription and coding.

The Inhibitory Control Battery Since measures obtained from single tasks can be noisy and contaminated by task-specific properties, we used four different inhibitory control tasks which we briefly describe below. **Go-NoGo task.** Participants were instructed to press a button as quickly as possible on each trial, except when a certain picture (e.g., a coconut) appeared (25% of the trails). One condition consisted of semantically-related objects (four fruits; 120 trials; 30 NoGo-Related), and the other of four semantically-

unrelated objects (120 trials; 30 NoGo-unrelated). **Fish-Flanker task** (Nozari, Trueswell, & Thompson-Schill, 2016). This was a version of the Flanker task with fish stimuli facing left or right (100 congruent, 100 incongruent trials). Embedded was a NoGo task in which the flanking fish had spots on them and the response had to be withheld (100 trials, NoGo-Fish). **Simon task**. A canonical Simon task with 80 congruent and 80 incongruent trials. **Picture Stroop**. Adopted from Nozari et al. (2016), this task was administered in four blocks. Each block contained only two pictures (e.g., pig/fox). In the first half of the block (8 trials), one of the two pictures appeared on each trial and participants named it. In the second half (8 trials) when each picture was presented, participants had to name the other picture. This situation produced a Stroop-like effect whereby the predominant response (picture's canonical name) had to be suppressed.

Response Coding/Scoring All responses were transcribed by two independent raters. Agreement between raters was 89% and discrepant cases were reconciled. Any complete or incomplete production of a non-target word, with or without correction, was coded as an error. For example “The eleph...no the snake next to the purple elephants are I mean is green.” was coded as containing an error on the target noun and an error on the verb. In addition, the presence or absence of disfluency was marked in four regions in the sentence: on or before the target noun, on or before the cue NP, on or before the verb, and finally, on or before the target cue. Disfluency was defined as unusually long pauses or prolongations that disrupted the flow of speech, filled pauses (uh, um, etc.), or repetitions (repairs are not coded, because they only happen on the primary “error” category which is analyzed separately). Not surprisingly, inter-rater agreement was lower for disfluency coding (69%), because of the subjective nature of coding silent pauses and prolongations which do or do not qualify as disfluency. To err on the side of caution, we excluded cases where the two raters did not agree on disfluent production.

For the inhibitory control battery, we used the effects on errors as our dependent measure of interest (agreement attraction), which also used a binary (error + disfluency) coding. For the NoGo tasks, the number of NoGo errors was used. For Flanker, Simon, and picture Stroop, the number of errors in the incongruent minus the number of errors in the congruent condition was used.

Results

Figure 2 shows the number of verb errors (upper panel) and the number of verb errors + disfluencies (lower panel) for each target-cue combination. We obtained 273 errors on verbs ($M = 2.2$, $SE = 0.39$ in the Target-flash condition; $M = 2.8$, $SE = 0.44$ in the Cue-flash condition). When disfluencies were added, this number increased to 487 verbs ($M = 3.72$, $SE = 0.55$ in the Target-flash condition; $M = 5.30$, $SE = 0.77$ in the Cue-flash condition).

Group-level Analyses

Group-level analyses were conducted using the package *LmerTest* in R v3.4.0. We used logistic multi-level models with random effects of subjects and items. We aimed to include a full random effect structure whenever possible, unless the full model did not converge. Since the general patterns of errors and disfluencies were similar (and so were the basic findings regarding attraction patterns on the two), we report the analyses on a combined set of errors + disfluencies which has a higher statistical power than errors alone. The first model tested for the general findings of agreement attraction reported previously in preamble tasks. The model contained the variables Attraction (SP/PS vs. SS/PP), Verb (is vs. are), and the interaction between the two, as well as random intercepts for subjects and items, and slope of attraction over subjects. This model revealed two critical findings: 1) There was a main effect of Attraction, such that there were significantly more errors + disfluencies on SP/PS trials than on SS/PP trials ($z = -3.31$, $p = 0.001$). 2) There was an interaction between Attraction and Verb, such that speakers made significantly more errors + disfluencies on SP (compared to SS) trials than they did on PS (compared to PP) trials ($z = -5.09$, $p < 0.001$). This shows the classic asymmetry in attraction errors reported by Bock & Miller (1991).

Next, we turned to testing how manipulating the inhibitory control demand affected agreement errors. This model included Attraction, Verb, Condition (Target-flash vs. Cue-flash), 2-way and 3-way interactions between these variables, as well as random intercepts for subjects and items. Due to space limitations, we focus on the theoretically important findings. The model confirmed the findings above: both the main effect of Attraction ($z = -3.94$, $p < 0.001$) and its interaction with Verb ($z = -2.04$, $p = 0.04$) were significant. In addition, there was a significant main effect of Condition, such that speakers produced significantly more errors + disfluencies in the Cue-flash than the Target-flash condition ($z = -2.19$, $p = 0.04$). Importantly, there was a significant 3-way interaction between Attraction, Verb, and Condition, such that the asymmetry in attraction was stronger in the Cue-flash condition ($z = 3.79$, $p < 0.001$). To unpack this interaction, we used two post-hoc models to test whether errors + disfluencies were significantly more common on SP (compared to SS) trials, PS (compared to PP) trials or both—perhaps with different magnitudes—in the Cue-flash compared to the Target-flash conditions. The first analysis was carried out on the subset of the data containing SP and SS trials, while the second analysis was conducted on a subset of the data containing PS and PP trials. The models included main effects of Attraction, Condition, and the interaction between the two, as well as random intercepts for subjects and items, and random slope of Condition over subjects. In the first model, all three effects were statistically significant (main effect of Attraction: $z = -6.90$, $p < 0.001$; main effect of Condition: $z = -2.10$, $p = 0.036$; interaction of

Attraction and Condition: $z = 3.83$, $p < 0.001$). In the second model, there was a marginal main effect of Attraction ($z = 1.96$, $p = 0.05$), but neither the main effect of Condition nor the interaction between Attraction and Condition reached significance. The results of these post-hoc tests indicate that the inhibitory control manipulation selectively increased the production of errors + disfluencies in the SP condition.




Figure 2: Mean proportion of agreement errors (upper panel) and errors + disfluencies (lower panel).

Analysis of Individual Differences

Analysis of individual differences was carried out by partial least squares path modeling, (PLS-PM) using the `plsmpm` package in R (Sanchez, 2013). PLS-PM is a partial least square approach to structural equation modeling suitable for analyzing the relationship between latent variables (psychological constructs, e.g., inhibitory control, agreement attraction) and manifest variables (observed data from the tasks we assume to measure these latent variables, e.g., Go-NoGo errors, agreement errors). One of the advantages of this kind of model over simple regression is that it allows for the inclusion of many manifest variables to represent the same latent variable without running into the problem of multicollinearity. This, in turn, allows for a much more robust and task-independent estimation of latent variables.

Recall that the results of the group analysis indicated that production of correct agreement under SP (but not PS) conditions required cognitive control. We thus constructed separate path models for SP and PS errors (each relative to its respective baseline SS and PP). If what the group analysis has really tapped into is inhibitory control, we would expect agreement errors + disfluencies in the SP (but not the PS) model to be predicted by the latent variable

inhibitory control. Figure 3 shows the architecture of the model which tests the contribution of inhibitory control (represented by the scores in our inhibitory control battery) to agreement attraction in SP trials (represented by errors + disfluencies on the SP minus the SS trials for Target-flash and Cue-flash conditions). A step-by-step tutorial on how to build and evaluate a PLS-PM has been provided in Nozari and Faroqi-Shah (2017). Here, we only highlight the key aspects directly relevant to the interpretation of the model's output. Both latent variables have high composite reliability values (0.85 for agreement attraction and 0.72 for inhibitory control), showing that they are well represented by their manifest variables. Also, the factor loadings (numbers outside of the parentheses on the connections between the latent and manifest variables) all have the same sign, showing that the effect of the latent variable on all the manifest variables is in the same direction (unidimensionality), which further shows that the latent variables are coherently represented by their manifest variables. Inspection of the factor loadings shows that some manifest variables have higher loadings on inhibitory control than others: the largest values belong to NoGo tasks, with smaller values for Flanker, Picture Stroop, and Simon tasks. In this model, the latent variable inhibitory control explains 20% of the variance on agreement attraction, and the connection between the two latent variables has a path coefficient of 0.45.




Figure 3: Structure of the PLS-PM testing the contribution of inhibitory control to agreement attraction in SP.

Significance testing was carried out via bootstrapping with 1000 iterations. The bold numbers in the parentheses show t -values. Not surprisingly, the manifest variables with high factor loadings all have t -values above 2. On the other hand, Flanker, Picture Stroop and Simon which had low factor loading on the inhibitory control variable, all have t -values equal or lower than 1, showing the greater contribution of NoGo scores. Most importantly, the t -value corresponding to the path coefficient between the latent variable inhibitory control and agreement attraction is 2.4, corresponding to a p -value of 0.017. This finding shows that the variability in production agreement attraction errors on SP trials can be explained by the variability in speakers' inhibitory control abilities. A PS model with an identical architecture, on the other hand, revealed no significant

effect of the latent variable of inhibitory control over agreement attraction errors produced in the PS minus PP conditions.

Discussion

We set out to test the potential contribution of inhibitory control to agreement production in native adult speakers of English. Results of group analysis using direct manipulation of inhibitory control demands and an analysis of individual differences using path modeling converged: production of the correct agreement in cases where a plural local noun competed with a singular subject was dependent on inhibitory control, and was predicted by individuals' performance on general inhibitory control tasks. We can thus conclude with certainty that agreement production, and more generally syntactic production, is not completely automatic.

Interestingly, not all inhibitory control scores were equally good predictors of agreement attraction problems in the analysis of individual differences. The effect was clearly stronger for the NoGo scores. One possibility is that NoGo scores have higher internal consistency compared to Flanker, Simon, and Stroop scores, all of which are subtraction scores with lower internal consistency. A more theoretically interesting possibility is that agreement production relies on a specific type of inhibitory control best indexed by NoGo scores. It has been proposed that the primary function of NoGo pathways is to delay, not suppress, actions (Munakata et al., 2011). Future empirical work should focus on pinpointing the exact type of inhibitory control underlying agreement production, and possibly other syntactic operations involved in sentence production.

In conclusion, this study took the first step in providing conclusive evidence for the relevance of inhibitory control to syntactic production. This finding calls for the inclusion of inhibitory control mechanisms in models of grammatical encoding in sentence production, examples of which we have proposed in Table 1. Moreover, it invites attention to deficits of non-linguistic control operations as a potential source for syntactic impairment in language disorders such as agrammatic aphasia.

Acknowledgments

This research was supported by the NSF grant 1631993 to Nazbanou Nozari and Akira Omaki.

References

- Badecker, W., & Kuminiak, F. (2007). Morphology, agreement and working memory retrieval in sentence production: Evidence from gender and case in Slovak. *Journal of Memory and Language, 56*(1), 65–85.
- Bock, K., & Miller, C. A. (1991). Broken agreement. *Cognitive Psychology, 23*(1), 45–93.
- Eberhard, K. M., Cutting, J. C., & Bock, K. (2005). Making syntax of sense: number agreement in sentence production. *Psychological Review, 112*(3), 531.
- Franck, J., Vigliocco, G., & Nicol, J. (2002). Subject-verb agreement errors in French and English: The role of syntactic hierarchy. *Language and cognitive processes, 17*(4), 371-404.
- Gleitman, L. R., January, D., Nappa, R., & Trueswell, J. C. (2007). On the give and take between event apprehension and utterance formulation. *Journal of Memory and Language, 57*(4), 544–569.
- Levelt, W. J. (1989). *Speaking: From intention to articulation*. ACL. MIT Press Series in Natural-Language Processing. MIT Press, Cambridge, Massachusetts.
- Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., & O'Reilly, R. C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Sciences, 15*, 453-459.
- Nozari, N., & Faroqi-Shah, Y. (2017). Investigating the origin of nonfluency in aphasia: A path modeling approach to neuropsychology. *Cortex, 95*, 119-135.
- Nozari, N., Freund, M., Breining, B., Rapp, B., & Gordon, B. (2016). Cognitive control during selection and repair in word production. *Language, cognition and neuroscience, 31*(7), 886-903.
- Nozari, N., Trueswell, J. C., & Thompson-Schill, S. L. (2016). The interplay of local attraction, context and domain-general cognitive control in activation and suppression of semantic distractors during sentence comprehension. *Psychonomic bulletin & review, 23*(6), 1942-1953.
- Sanchez, G. (2013). PLS path Modeling with R Trowchez editions. Berkeley. <http://www.gastonsanchez.com/PLS Path Modeling with R.pdf>.
- Shao, Z., Meyer, A. S., & Roelofs, A. (2013). Selective and nonselective inhibition of competitors in picture naming. *Memory & Cognition, 41*(8), 1200–11.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84*(2), 127.
- Trude, A., & Nozari, N. (2017). Inhibitory control supports referential context use in language production and comprehension. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society (pp. 1218-1223). Austin, TX: Cognitive Science Society.