

# Communicatively efficient language production and case-marker omission in Japanese

Chigusa Kurumada

kurumada@stanford.edu

Department of Linguistics

Stanford University

T. Florian Jaeger

fjaeger@bcs.rochester.edu

Department of Brain and Cognitive Sciences

University of Rochester

## Abstract

Recent studies hypothesize that language production is governed by the principle of efficient information transmission: Speakers tend to omit elements whose information content is contextually predictable, while providing more linguistic signal to convey otherwise less predictable information. However, previous findings in support of this hypothesis are also compatible with alternative accounts based on production difficulty. To distinguish between these competing accounts, we conducted experiments on speaker's preference in optional case-marking in Japanese. The results suggest that Japanese speakers are more likely to omit the object case-marker when an associated noun has properties (e.g., animacy) that are prototypical to a grammatical object. Moreover, case-marker omission was facilitated when other elements in a sentence made the grammatical function assignment more predictable. The results were obtained with all the factors related to production difficulty held constant, and thus provide support for the models of communicatively efficient language production.

**Keywords:** Case-marker omission, Japanese, language production, efficiency, information transfer

## Introduction

Language often offers multiple options for expressing semantically equivalent or near-equivalent messages. For example, speakers have a choice between different word orders and voice choices (e.g. the ditransitive alternation in English; active vs. passive; scrambling in languages like German, Japanese, and Hindi). Another common type of alternation involves the choice between more or less linguistic form (e.g., phoneme duration, morphemes, phonological alternation) to encode the same meaning. For example, in English sentence like (1), speakers can, but do not have to, produce the complementizer “that”. The past psycholinguistic studies have used these alternations as windows into the cognitive processes underlying language production.

(1) My boss said (that) we were absolutely crazy.

Recent psycholinguistics work has proposed that speakers' preferences in such alternations provide evidence that the computational system underlying language production is organized to facilitate efficient information transfer (e.g., Aylett & Turk, 2004; Levy & Jaeger, 2007; Jaeger, 2010). This line of work has focused on the observation that it seems to be predictable linguistic material, and hence material that is low in Shannon information, that tends to be reduced or even completely omitted (e.g., Bell, Brenier, Gregory, Girand, & Jurafsky, 2009; Frank & Jaeger, 2008; Resnik, 1996). For example, whether English speakers produce the complementizer *that* depends on how predictable it is for the continuation

of words so far (e.g., My boss said) to take a complement clause, as in (1): when the complement clause is much expected, *that* is likely to be omitted; when it is less so, *that* is likely to be inserted (Jaeger, 2010).

Observations like these can be accounted for in terms of communicative efficiency since the reduction/omission of predictable material and lengthening/insertion of less predictable material results in more uniform distribution of information density across the speech signal, which is proven to facilitate efficient -i.e. fast and robust- information transfer (Genzel & Charniak, 2002) and to minimize processing difficulty (Levy & Jaeger, 2007). A trade-off between the amount of information and the amount of linguistic signal expended is also expected under models of boundedly rational communication: If speakers are rational and aim to balance production effort with communicative success, they should provide a more perceptual signal when a word or structure is less contextually inferable (the Ideal Speaker Model in Jaeger, 2011, see also Piantadosi, Tily, & Gibson, 2011).

However, for much of the evidence that has been cited in favor of communicative efficiency accounts, it is an open question to what extent it could be accommodated in competing accounts. In particular, it is well-established that difficulty with the retrieval, processing, or articulation of upcoming material can affect the degree of reduction/omission of (typically immediately preceding) material. This includes effects on phonetic reduction (Fox Tree & Clark, 1997) and the optional production of “that” (e.g., Ferreira & Dell, 2000) described above. For example, speakers are more likely to produce the optional “that”, if the onset of the complement clause (its subject) is frequent, short and has previously been mentioned (Ferreira & Dell, 2000; Roland, Elman, & Ferreira, 2005). We follow previous work and refer to these effects as availability-based effects.

Here we investigate a morpho-syntactic alternation that provides a less ambiguous test case for the communicative-efficiency accounts. We present three production experiments on optional case-marking in Japanese. Japanese is a verb-final language with relatively flexible word order: in a transitive sentence like (2), both SOV and OSV are possible word orders, although the latter is considerably less frequent. Case relations are marked with post-nominal particles (case-markers) as shown in (2), where “-ga” is the nominative marker and “-o” is the accusative marker. Here we focus on the optionality of the *accusative* case-marker, indicated in () by parentheses.

(2) Taro-ga sushi-(o) tabe-ta.  
 Taro-NOM sushi-(ACC) eat-PAST.  
 Taro ate sushi.

Compared to languages like English or Mandarin Chinese, the flexibility in Japanese word order implies higher uncertainty about the grammatical function assignment (henceforth GF-assignment). In other words, it is not immediately obvious what is the subject and what is the object especially when a case-marker is not present. While case-markers are obligatory in written discourse, however, they are frequently omitted in conversational speech — often, like in (2), without change in meaning.

Such optional case-marking has received considerable attention in linguistic work (e.g., Aissen, 2003), but has remained comparatively understudied in psycholinguistics. Here it is of interest because it allows us to distinguish between the predictions of communicative efficiency accounts, specifically the ideal speaker model proposed in Jaeger (2011), and those of availability-based accounts.

The ideal speaker model predicts that speakers should be more likely to produce case-marking when the intended GF-assignment is unexpected given the other information provided in the sentence (for details, see Jaeger, 2011). This is the prediction we test here. In the remainder of this paper, we will use the more succinct statement that “speakers should be more likely to mark the unexpected” to refer to this prediction. In Experiment 1, we investigate whether object typicality, manipulated via changes in animacy, affects speakers’ preference during production of optional case-marking in Japanese. We compare the production of sentences that only differ in the animacy of the object, which is either a human referent (e.g., the student) or an inanimate referent (e.g., the fire-engine)(Figure 1-a). If speakers are more likely to mark the unexpected, they should be more likely to produce object case markers if the object is human [atypical] compared to if it is inanimate [typical] (Aissen, 2003).

Preliminary evidence for this prediction comes from qualitative work (Lee, 2007), which examined optional case-marking in conversational Korean. She tested the hypothesis suggested by typological work on case-marking systems. Lee found that both definiteness and animacy are significant predictors of Korean speakers’ use of subject- as well as object case-markers in the direction predicted here. However, in a similar corpus study in Japanese, Fry (2003) found no effect of animacy in the object-case marking. While inanimate and indefinite (i.e. atypical) subjects were significantly more likely to be case-marked than animate and definite subjects (69% vs. 64%), he found no effect of animacy on object case-marking. It is unclear, however, whether this result indicates an actual absence of an effect of animacy, or whether this was simply due to the sparseness of data or confounding factors, not controlled in the corpus study. Experiment 1 addresses this question by manipulating *only* object animacy. Further, items were constructed such that all materials following the direct object were held constant, so that no differ-

ences in case-marking preferences between the conditions are predicted by availability-based production (Ferreira & Dell, 2000).

Experiment 2 and 3 put the ideal speaker model’s prediction to a stronger test. We investigate if Japanese speakers’ case-marking preference can be directly affected by the inferability of GF-assignment, *beyond categorical factors like animacy, givenness, or definiteness*. To this end we manipulate *plausibility* of GF-assignments and investigate their effect on the production of case markers.

## Experiment 1

We employ a spoken recall paradigm to test if Japanese speakers’ use of the accusative case-marker *-o* is sensitive to animacy of the direct object. We manipulated the animacy of the direct object and whether the original stimulus contained a direct object case-marker or not. If optional case-marking is affected by a preference for communicative efficiency, speakers should be more likely to produce responses with a case-marker for animate (atypical) objects compared to inanimate (typical) objects.

### Methods

**Participants** 20 native speakers of Japanese in Stanford area participated in this study. They received \$7 for their participation.

**Materials** In a Latin-square design, each list contained 24 items and 48 fillers. As described above, items were transitive sentences with either animate or inanimate direct objects and with or without a case-marker. Sentence patterns for recall stimuli are illustrated in Figure 1. The nominative case-marker was always present, avoiding ambiguity about the intended GF-assignment and hence about the meaning. Additionally, all items were presented in the subject-before-object order, which is hugely more frequent in Japanese.

Fillers were length-matched sentences with intransitive verbs and longer adverbial phrases. There was no lexical overlap between any of the stimuli. Stimuli were grouped into pairs so that there were 24 item-filler pairs and 12 filler-filler pairs, totalling 36 trials. The order of items and fillers within a pair and the order of pairs were held constant across participants. All stimuli were recorded by a female native speaker of Japanese, using the same prosody for all conditions. In addition, the same speaker recorded 72 recall prompts, one for each sentence (always the verb).

**Procedure** Each trial consisted of an encoding phase and a recall phase. During encoding, participants listened to pairs of sentences and were instructed to remember them. During recall, participants heard the verb of one of the two sentences (the prompt). They then recalled and produced the full sentence corresponding to that verb. Subsequently, the second prompt was played and participants produced the second sentence. In half of the trials, the sentence encoded first was also recalled first. In the other half, the order was reversed. Following standard procedure (Ferreira & Dell, 2000), target

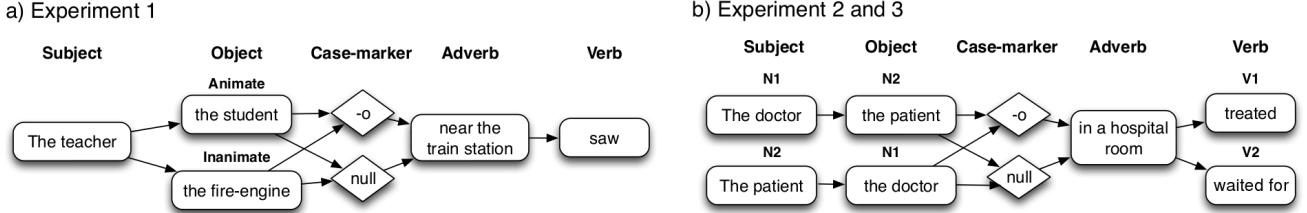


Figure 1: Sentence patterns in a) Experiment 1 (left) and b) Experiment 2 and 3 (right).

Table 1: Percentage of recall error for each of the four conditions in Experiment 1.

	Animate	Inanimate
Present	8.9%	8.9%
Absent	15.6%	8.9%

items were never recalled directly after encoding. That is, if an item was encoded first, it was recalled either first or second. If it was encoded second, it was always recalled second. Across participants the recall order for each pair of stimuli was held constant.

**Scoring** All  $20 \times 24 = 480$  recorded items were transcribed and coded by a native Japanese undergraduate research assistant who was unaware of the purpose of this experiment. Sentences with recall errors anywhere in the sentence were excluded (11.6%), leaving 424 responses for our analysis. The error rate was doubled for sentences with animate objects without case-marker compared to the other three conditions (Table 2,  $\chi^2(3) = .7$ ,  $p < .05$ ). Although these sentences are technically not ambiguous (recall that all sentences contained subject case-marking), sentences with animate objects that lack case-marking are the ones we hypothesized to be most confusing. The high error rate indicates that the recall production was in fact difficult with this type of sentences.

We also coded the presence of disfluency (88 out of the 424 analyzable sentences contained at least one disfluency). There was no significant difference across sentence types ( $\chi^2(1) = .03$ ,  $p > .8$ ).

## Results and discussion

A mixed logit regression analysis with the full 2 (animacy of object)  $\times$  2 (presence of case-marker in stimulus) design was employed to predict the presence of a case-marker. We report the results for the model with maximum random effect structure justified by the data based on model comparison (Jaeger, 2008), which contained random by-subject and item intercepts as well as by-item slopes for the presence of case-marker in stimulus. The effect of interest reported below was robustly significant even in the full random effect model (by-subject and by-item slopes for all factors and their interaction). All predictors were contrast coded and there were no signs of collinearity (fixed effect  $rs < .07$ ).

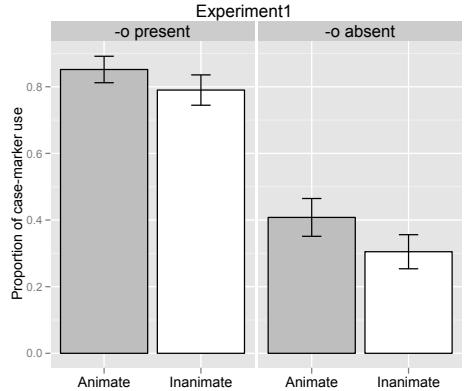


Figure 2: Proportion of case-marker use in Experiment 1. The error bars represent  $\pm 1$  standard error.

Unsurprisingly, participants were more likely to produce a object case-marker if the original stimulus contained a object case-marker ( $\beta = 2.3$ ,  $p < .0001$ ). Crucially, participants were also more likely to produce a case-marker if the object was animate ( $\beta = .45$ ,  $p < .03$ ). The two factors did not interact ( $p > .9$ ). Figure 2 summarizes the two main effects.

In this recall production experiment, we thus found that animacy of direct objects affected Japanese speakers' case-marking preferences. This suggest that the failure to find animacy-based effects on object case-marking in casual speech in previous work is indeed due to a lack of power and confounding effects of other variables that optional case-marking is sensitive to (Fry, 2003). Our results hence replicate for Japanese what has been found in corpus-based studied on Korean speech (Lee, 2007). Our findings are also compatible with the qualitative description of differential case-marking across languages of the world (Aissen, 2003).

Most relevant to the current purpose, the animacy-based effect on case-marking support the ideal speaker model. The effects we observe are expected under the ideal speaker model and cannot be reduced to availability-based production since the material following the direct object was held constant within items. As discussed above, it is also not clear how production-oriented ('accessibility-based') accounts of the type that have been proposed for phonetic reduction (Arnold, in press; Bell et al., 2009) would account for the observed effects. More generally, it seems unlikely that the effects

observed here are reducible to production difficulty. While we found an increase proportion of recall errors in sentence with animate (atypical) objects without case marker, the proportion of recall errors for animate objects *with* case markers was identical to that of inanimate objects. Furthermore, the interaction observed for recall errors is not observed in our analysis of speakers' case-marker preferences. Finally, as mentioned above, we did not observe any differences in the distribution of disfluencies across the animacy conditions.

## Experiment 2

If the animacy effect we saw in Experiment 1 was indeed due to effect of communicative efficiency on the encoding of GF-assignment, the effect should remain when the animacy of the arguments are controlled. Using animate-animate noun pairs, we examined if the plausibility of GF-assignment affects Japanese speakers' use of the accusative case-marker.

In Figure 1-b (right), the GF-assignment is more plausible when the doctor treats the patient rather than the other way around. Hence the ideal speaker model predicts that speakers should be more likely to produce the object case-marker when the doctor is the object of the sentence, compared to when the patient is the object.

## Methods

**Participants** 32 native speakers of Japanese in Stanford area participated in this study. They received \$7 for their participation.

**Stimuli** There were 24 items, consisting of subject, object, adverb, and verb. Each item consisted of eight conditions (=192 stimuli), resulting from crossing 1) the plausibility of GF-assignment based on the order of two noun phrases and the verb, 2) presence/absence of the case-marker, and 3) the identity of the verb. This is illustrated in Figure 1.

In the example in Figure 1-b, plausibility was coded as *high* in "The doctor treated the patient in a hospital room" and "The patient waited for the doctor in a hospital room" and as *low* in "The patient treated the doctor in a hospital room" and "The doctor waited for the patient in a hospital room". Below we collapse over the 2-way within-item verb contrast and treat the design as a 2x2 (plausibility by case-marker presence in the input), since the identity of the verb is of no theoretical interest here (the effects of verbs on plausibility are already capture by our coding of plausibility). Furthermore, the verb identity was balanced within item and within the two plausibility conditions.

As in Experiment 1, the 24 items were combined with 48 length-matched fillers in a Latin-square design that held order of stimuli constant across lists. There was no lexical overlap between any of the stimuli. The grouping of stimuli into pairs was the same as in Experiment 1. All stimuli were recordings of a female native speaker of Japanese. In addition, the same speaker recorded 72 prompts, one for each sentence (always the verb).

Table 2: Percentage of recall error for each of the four conditions in Experiment 2.

	Plausible object	Implausible object
Present	8.6%	16.7%
Absent	8.6%	12.3%

**Norming studies** The stimuli described below were created using an online norming study (40 native speakers of Japanese). Participants were presented the 24 animate-animate noun pairs used in our items (e.g., doctor-patient). Participants first rated the relative naturalness of the two potential patterns of grammatical function assignment on a 10 point rating scale (e.g., *The doctor (did something to) the patient* vs. *The patient (did something to) the doctor*). Second, we asked them to provide verbs that would make the most natural continuation of a given pair of nouns for each pattern (e.g. doctor-NOM patient-ACC / patient-NOM doctor-ACC). Among the set of verbs given by the informants, we selected two verbs for each noun pair such that one verb maximizes the object probability of one noun, and the other verb maximizes the object probability of the other noun. For example, for the "doctor-patient" pair, we chose "(to) treat", which makes "patient" as the plausible object, and "(to) wait for", which makes "the doctor" more plausible.

A second norming study (40 native speakers of Japanese; no overlap in participants with first norming study or either of the experiments) asked participants to rate the relative naturalness of the two patterns of GF-assignment when the verb is present (e.g., "doctor-NOM patient-ACC hospital-room-LOC treat" vs. "patient-NOM doctor-ACC hospital-room-LOC treat").

**Procedure** The procedure was identical to that of Experiment 1. Since some of the stimuli (e.g., *The patient treated the doctor*) were meant to be less plausible, participants were instructed to listen to sentences carefully and produce them faithfully to the input even when they were "somewhat surprising".

**Scoring** All  $32 \times 24 = 744$  responses were transcribed and coded by a native speaker of Japanese. Sentences with recall errors were excluded (11.6% [sic]), leaving 658 responses for our analysis. Numerically, the error rate was higher for sentences with implausible GF-assignments (14.5% vs. 8.6%) but the difference was not statistically significant ( $\chi^2(1) = .24$ ,  $p > .62$ ). 13% of the error free sentences contained at least one instance of disfluency, but the occurrence rate did not differ across conditions ( $\chi^2(3) = .80$ ,  $p > .3$ ). As in Experiment 1, the distribution of disfluencies did not differ across conditions ( $\chi^2(3) = .80$ ,  $p > .3$ ).

## Results and Discussion

Using the same statistical approach as in Experiment 1, we analyzed the remaining 2x2 design defined by the plausibility of GF-assignment (high vs. low) and case-marker presence

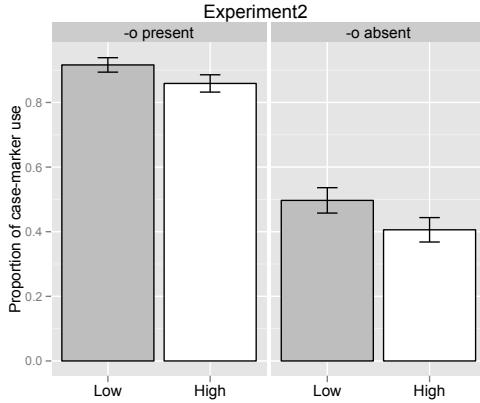


Figure 3: Proportion of case-marker use in Experiment 3 by plausibility of GF-assignment (high vs. low) and presence of the object case marker in the recall stimulus (-o present vs. absent). The error bars indicate  $+/- 1$  standard error.

in the originally presented stimuli (presence vs. absence) to predict the presence of a case-marker in the responses. Unsurprisingly, participants were more likely to produce an object case-marker if the original stimulus contained an object case-marker ( $\beta = 2.6$ ,  $p < .0001$ ). More importantly, they were also more likely to produce a case-marker for implausible GF-assignment ( $\beta = 0.8$ ,  $p < .003$ ). The two factors did not interact ( $p < .7$ ). Thus, this model confirmed the predicted effect of object plausibility in the form of categorical factor. The effects are illustrated in Figure 3.

Next, we analyzed the gradient effects of values for the plausibility of GF-assignments as rated in the norming studies described above. The effect of two different plausibility ratings was examined: plausibility of GF-assignment given only the two nouns (Norming Study 1, part 1) or given the full sentence (Norming Study 2). Plausibility ratings for full sentences returned the expected significant effect ( $\beta = -0.43$ ,  $p < .03$ ), namely that the speakers were more likely to use the case-marker when the GF-assignment had been normed to be less plausible. However, the ratings based on only the two noun phrases returned no such effect ( $\beta = -.05$ ,  $p > .9$ ). This suggests that what speakers assess is the plausibility of an intended message (including the verb information) rather than just the properties of nouns.

### Experiment 3

We replicated Experiment 2 by using the subject noun, instead of the verb, as a recall cue. This allows us to rule out the possibility that the plausibility effect observed in Experiment 2 is an artifact of the recall cue being the verb of a sentence. As Japanese is a SOV language, it is possible that the use of the (sentence final) verb as the recall cue in Experiment 1 and 2 forced participants to adopt task-specific production strategies that do not reflect normal language production.

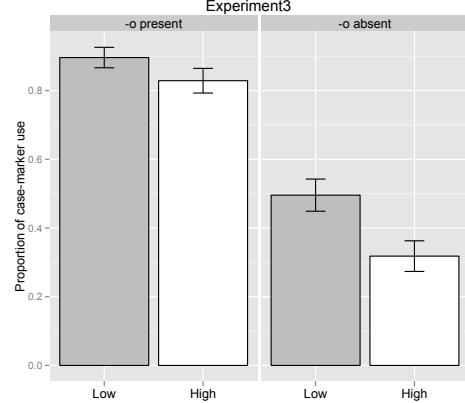


Figure 4: Proportion of case-marker use in Experiment 3 by plausibility of GF-assignment (high vs. low) and presence of the object case marker in the recall stimulus (-o present vs. absent). The error bars indicate  $+/- 1$  standard error.

## Methods

**Participants** 26 native speakers of Japanese in Stanford area participated in this study. Data from one subject were excluded because of a problem with the recording device.

**Stimuli** The stimuli and their presentation order were mostly identical to those of Experiment 2. New recall cues (the subject nouns of the target sentences) were recorded by the same speaker who recorded the stimuli sentences for Experiment 2.

## Results and Discussion

Using the same statistical approach as in Experiment 1 and 2, we analysed the binary predictors of the plausibility of GF-assignment (high vs. low) and case-marker presence in the originally presented stimuli (presence vs. absence of -o) to predict the occurrence of a case-marker in the responses. We replicated the finding in Experiment 2: the expected effect of the object case-marker presence in the original stimulus ( $\beta = 1.83$ ,  $p < .001$ ), and the low inferrability GF-assignments ( $\beta = -1.12$ ,  $p < .007$ ). The two effects did not interact ( $p > .2$ ). Thus, this model confirmed that the effect of object plausibility (in the form of categorical factor) remains when the recall cue was the subject of the sentence. The effects are illustrated in Figure 4.

Also, we replicated the effects of two different plausibility ratings. Plausibility ratings for full sentences returned the expected significant effect ( $\beta = -0.40$ ,  $p < .05$ ) whereas the ratings based on only the two noun phrases returned no such effect ( $\beta = -.03$ ,  $p > .6$ ). There was a negative interaction term between the plausibility value given the full sentence and the presence of the case-marker: Highly plausible GF-assignments were even less likely to be case-marked when the original sentences lacked the object case-marking ( $\beta = -.74$ ,  $p < .01$ ).

## General Discussion

Our results suggest that Japanese speakers prefer to produce an object NP without case marking when grammatical function of a noun is made more predictable given the semantics of the noun (e.g., animacy) and the other linguistic elements in the sentence (e.g., plausibility of GF-assignment given the subject, object, and verb). The plausibility effect we saw in Experiment 2 and 3 strongly suggests that speakers have fine-grained probabilistic knowledge about the plausibility beyond categorical factors like animacy and definiteness. More generally, this is one of the first studies showing a systematic effect of plausibility in the morpho-syntactic encoding of speech.

Recent studies have found that speakers' use and non-use of case-markers are interacting with the word order as well. Speakers tend to use case-makers when the word order does not conform to an expected pattern (see Lee and Kim (2012) in Korean, Fedzechkina, Jaeger, and Newport (2012) for experiments using an artificial language). This is also predictable by the assumption of efficient communication. An explicit case marking becomes more likely when the non-canonical word order biases against an intended GF-assignment pattern.

This body of research including the current study constitutes strong support for the view that language production is optimized to maximize the efficiency of information transmission (Jaeger, 2010; Levy & Jaeger, 2007), and, in particular, the ideal speaker model (Jaeger, 2011). Unlike most previous work, the current results cannot be accounted for in terms of availability, even if availability-based production is extended to include the availability of upcoming syntactic structures.

Furthermore, our results provide broader implications for studies investigating the effects of communicative pressure on cross-linguistically attested phenomena such as differential case-marking (Aissen, 2003). Our results and Lee's (2007) study in Korean support the hypothesis that the languages with an optional case-marking system are sensitive to the same factors that are known to affect more categorical case-encodings in various languages. This may suggest that the functional pressure for efficient communication underlies at least some of the universal features found across languages.

## References

Aissen, J. (2003). Differential object marking: iconicity vs. economy. *Natural Language and Linguistic Theory*, 21, 435–483.

Arnold, J. E. (in press). Reference production: Production-internal and addressee-oriented processes. *Language and Cognitive Processes*.

Aylett, M. P., & Turk, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and Speech*, 47(1), 31–56.

Bell, A., Brenier, J., Gregory, M., Girand, C., & Jurafsky, D. (2009). Predictability effects on durations of content and function words in conversational English. *JML*, 60(1), 92–111.

Fedzechkina, M., Jaeger, T. F., & Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*.

Ferreira, V., & Dell. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology*, 40, 296–340.

Fox Tree, J. E., & Clark, H. H. (1997). Pronouncing “the” as “thee” to signal problems in speaking. *Cognition*, 62, 151–167.

Frank, A., & Jaeger, T. F. (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. In *Proceedings of the 30th annual meeting of the Cognitive Science Society*.

Fry, J. (2003). *Ellipsis and wa-marking in japanese conversation*. New York: Routledge.

Genzel, D., & Charniak, E. (2002). Entropy rate constancy in text. In *Proceedings of the Association of Computational Linguistics* (pp. 199–206). Philadelphia, PA.

Jaeger, T. F. (2008). Categorical data analysis: Away from anovas (transformation or not) and towards logit mixed models. *JML*, 59, 434–446.

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, 61, 23–62.

Jaeger, T. F. (2011). Communicative efficiency and adaptiveness in the ideal speaker (ms). *University of Rochester*.

Lee, H. (2007). Case ellipsis at the grammar/pragmatics interface: A formal analysis from a typological perspective. *Journal of Pragmatics*, 39(9), 1465 – 1481.

Lee, H., & Kim, N. (2012). Non-canonical word order and subject-object asymmetry in korean case ellipsis. In *Proceedings of the 19th international conference on head-driven phrase structure grammar* (pp. 427–442). Stanford, CA: CSLI Publications.

Levy, R., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. In B. Schrödlkopf, J. Platt, & T. Hoffman (Eds.), *Advances in Neural Information Processing Systems (NIPS)* (Vol. 19, p. 849–856). Cambridge, MA: MIT Press.

Piantadosi, S., Tily, H., & Gibson, E. (2011). Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Sciences*, 108(9), 3526.

Resnik, P. (1996). Selectional constraints: An information-theoretic model and its computational realization. *Cognition*, 61, 127–159.

Roland, D., Elman, J. L., & Ferreira, V. S. (2005). Why is that? Structural prediction and ambiguity resolution in a very large corpus of English sentences. *Cognition*, 1–28.