To appear in "Selected papers from DAARC 2011", to be published by Springer. Eds. A. Branco, I. Hendrickx, R. Mitkov, Sobha L.

Exploring the Lexical and Acoustic Consequences of Referential Predictability

Elsi Kaiser, David Cheng-Huan Li and Edward Holsinger

Department of Linguistics, University of Southern California
Los Angeles, CA 90089, USA
{emkaiser, lidc, holsinge} @ usc.edu

Abstract. Findings from various domains suggest that predictability is an important component of language processing. We report psycholinguistic research suggesting that predictability also influences referential processing, in the form of reduced acoustic durations for predictable referents. However, we do not find evidence that predictability directly influences likelihood of pronominalization, contrary to some prior claims. Instead, our data indicate that the use and interpretation of pronouns is influenced by thematic role, independently of which referent is most predictable, i.e., most likely to be mentioned next. We suggest that likelihood-of-mention is influenced by the mapping between syntactic and thematic roles. Our results highlight the benefits of exploring both lexical and acoustic aspects of referential production.

Keywords: psycholinguistics, reference resolution, anaphora, predictability, acoustic reduction

1 Introduction

A growing body of research indicates that the language processing system is sensitive to predictability. In the domain of phonology, for example, words are produced with shorter durations and greater phonological reduction when they are predictable on the basis of preceding words/collocational frequencies ([1], [2]). Predictability effects have also been explored in areas such as syntax ([3], [4]) and language change ([5]).

We report two experiments that investigate the role of predictability in reference-tracking, focusing especially on the relation between (i) how an entity is referred to and (ii) how likely it is to be mentioned, i.e., how predictable it is.

Given the importance of predictability in many models of human language processing, it may seem surprising that in the domain of reference-tracking, there is as of yet no consensus regarding the role of predictability. One view is that referents' predictability – likelihood of being mentioned subsequently – is closely connected to the choice of referring expression: Reduced forms (e.g. pronouns) are preferentially used for highly predictable, expected referents (e.g. [6], [7]).

In contrast, others claim that likelihood-of-mention is separate from choice of referring expression. For example, Fukumura & van Gompel [8] argue that likelihood-of-mention is influenced by semantic factors, and that referring expression

choice is influenced by accessibility/structural prominence, with pronouns preferring syntactically-prominent referents/subjects. In related work, Kehler, Kertz, Rohde & Elman's [9] Bayesian approach to pronoun comprehension and production makes it possible to separate likelihood-of-mention from likelihood-of-prominalization.

$$P(referent|pronoun) = \frac{P(pronoun|referent)\ P(referent)}{P(pronoun)}$$

Fig. 1. Kehler et al.'s Bayesian approach

According to this approach, comprehenders have expectations about who is likely to be mentioned next in a discourse. Kehler et al. suggest that when considering pronoun interpretation patterns, represented by (*P(referent|pronoun*) – given a pronoun, what is the probability that it refers to a particular referent? – it is useful to treat interpretation biases as being composed of two expectations: (i) the probability that a particular referent will be mentioned subsequently (*P(referent)*) (see also [10]) and (ii) the probability that a particular referent will be referred to with a pronoun (*P(pronoun|referent*)). Pronoun interpretation is of course also sensitive to the overall probability of using a pronoun, *P*(*pronoun*). Thus, from the perspective of the comprehender, pronoun resolution involves *P(referent|pronoun*), i.e., given a pronoun, what is the probability that it refers to a particular referent x?

Crucially, according to this approach, likelihood-of-mention (*P(referent)*) can be separated from likelihood of producing a pronoun (*P(pronoun|referent*)) and does not necessarily map directly to pronoun interpretation (*P(referent|pronoun)*).

Further evidence pointing to a dissociation between likelihood-of-mention and pronominalization comes from Kaiser's [11] data on *it*-clefts (e.g. *It was Mary that Kate tickled*). In an open-ended production task after subject clefts, participants were most likely to pronominalize reference to the immediately preceding subject (i.e., *P(pronoun|subject)* was relatively high), but this character was not overall the most likely referent to be mentioned next. Although we will not be framing the experiments in this paper in specifically Bayesian terms, the approach sketched out above helps to show how likelihood-of-mention and pronominalization can be separated. This differs from the Expectancy Hypothesis [6], which suggests that the expectancy /predictability of different entities is closely tied to speakers' choice of referring expression, with more attenuated forms used for more expected/predictable referents.

In sum, existing research on the relationship between likelihood-of-mention (how predictable, how expected a particular referent is at a certain point in the discourse, which can be formalized as *P(referent)* in Kehler's terms) and likelihood-of-pronominalization – both in terms of how likely speakers are to use pronouns (e.g. *P(pronoun|referent*)) and how likely hearers are to resolve pronouns in a particular way (*P(referent|pronoun)*) – has led to divergent claims.

2 Overview of Experiments

To shed light on the debate regarding the relation between predictability and referential form, we explored effects of likelihood-of-mention in two different dimensions: (i) on the production and choice of referring expression and (ii) the

acoustic duration of referring expressions. As a starting point, we investigated *pronoun interpretation:* When given a pronoun, how do people interpret it (Experiment 1)? Then, in Experiment 2 we investigated how pronoun interpretation compares to *pronoun production*, i.e., in a production task, when do people opt to use pronouns rather than names or other referring expressions? Crucially, we also probed which referents participants choose to talk about, thereby obtaining a measure of how predictable different referents are.

With this methodology, we obtain two kinds of information about pronominalization that allow us to probe the relation between pronominalization and likelihood-of-mention: (i) *Comprehension*: When faced with a pronoun, does a person tend to interpret it as referring to the most predictable referent? and (ii) *Production*: When producing pronouns, do people produce pronouns mostly when referring to the most predictable referent? In addition, the production task (Experiment 2) allows us to test for potential effects of acoustic reduction.

In this paper, we use the terms predictability and likelihood-of-mention to refer to both speakers and comprehenders. Comprehenders have probabilistic expectations about who will be mentioned next ([9], [10]), and thus, from the comprehender's perspective, referents vary in how predictable / expected they are. Crucially, following [12] and [9], we regard comprehenders as "savvy anticipators" who have forward-looking expectations about who the speaker will talk about next and with what kind referring expression (see Fig.1). These expectations are shaped by prior context and continuously updated over the course of a discourse on the basis of new linguistic input (see also [10]). Broadly speaking, according to this view, under normal circumstances the comprehender's expectations about which referents are most predictable should largely mirror the speaker's expectations (see Fig.1).

3 Experiment 1: Prompt pronoun

3.1 Participants, materials and design

Twenty-four native English speakers from the University of Southern California community participated. Participants heard auditorily-presented sentences (ex.1) with agent-patient verbs (e.g. AGENT *kicked* PATIENT, AGENT *tickled* PATIENT) over headphones, while viewing scenes depicting the two mentioned characters and other objects. The two characters always had the same gender (two women or two men). We used four male and four female characters throughout the study (female: *Lisa, Kate, Mary, Anne*; male: *Peter, Greg, John, Mike*). At the start of the experiment, participants completed a familiarization procedure to learn which name maps onto which character. However, because we did not want to induce a memory burden or induce potential misunderstandings, the name of each character was also included on the displays, as shown in Fig.2. On critical items, the two characters were located on the left and right sides of the display, with other objects between them (Fig.2). The left/right location of subjects and objects was counterbalanced. We also

counterbalanced how often each character appeared on the left and on the right, and in subject position and object position.

(1a) Mary slapped Lisa at the zoo. As a result she…
(1b) Lisa was slapped by Mary at the zoo. As a result she…

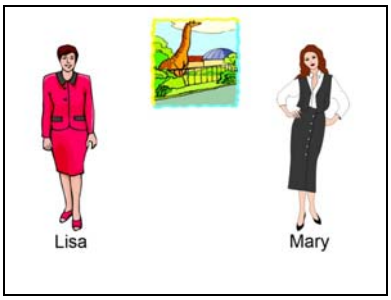

Fig. 2: Sample image

The critical sentences were in the active or the passive voice, and were followed by a sentence fragment consisting of a connective and a pronoun (ex.1). The participants' task was to produce natural-sounding continuations using the sentence fragment. This task can be regarded as a combination of comprehension and production, as participants need to understand the critical sentence and – crucially for our purposes – need to interpret the pronoun in the sentence fragment before they can provide a continuation (see [10], [12] for discussion of this task).

All critical sentences ended in a location expression (e.g. *at the museum, on the beach*) which was related to the objects between the two characters. This was included because participants' eye-movements were recorded during this task, and the location expression acted as a look-away, encouraging participants to look away from both characters before they hear the pronoun. (Eye-movements are not reported here.)

Because coherence relations between sentences influence discourse flow (e.g., [9], [13], [14]), we used the causal 'as a result' connective, to ensure consistency in the coherence relations that participants construct.[1] More specifically, causal connectives after agent-patient verbs are known to focus the patient thematic role (see e.g. [15]). Thus, we used 'as a result' to create a situation in which there is a bias towards the patient, at least in the active voice. The behavior of subject pronouns after a passive clause and a causal connective is not yet well-understood, and thus we wanted to test whether the predicted patient bias persists in passives. As a whole, the study included 24 targets (12 with 'as a result' (6 actives, 6 passives), and 12 with 'then' connectives (6 actives, 6 passives, see footnote 1) and 36 fillers. We used a Latin-Square design to construct four lists. Reverse lists were also constructed to control for trial order.

[1] Experiments 1 and 2 also included *then* connectives, which are ambiguous between a temporal interpretation (*Peter came home. Then he ate dinner*) and a causal interpretation (*Peter kicked Andy. Then Andy started to cry*). Preliminary analyses suggest that when used causally, *then* resembles *as a result*. However, due to the ambiguity of *then* and the resulting unbalanced data points, in this paper we focus on *as a result* only, for both experiments.

3.2 Coding

All continuations were double-coded by two coders working independently. Coders noted whether the prompt pronoun was used to refer to the preceding subject, the preceding object, or whether its antecedent was unclear. We use the term 'subject' to refer to the grammatical subject of the critical sentence, i.e., the agent in actives and the patient in passives. We use the term 'object' to refer to the direct object in actives (the semantic patient) and the object of the *by*-phrase (the semantic agent) in passives.

Each coder went through the data independently. Coders were instructed to be conservative and to avoid over-interpretation, i.e., to err on the side of choosing 'unclear' if there was not enough information available to determine the intended antecedent. Subsequently, any discrepancies between the coders were resolved through discussion. If the two coders did not agree on a referent (i.e. if a discrepancy could not be resolved), the pronoun was coded as 'unclear.' In the end, 25% of pronouns in the 'result' conditions that are reported here were coded as having unclear/ambiguous antecedents; these were excluded from subsequent analyses. This kind of coding procedure is similar to [9], [10], [11] and [12]. Examples of participants' continuations and how they were coded are provided in ex(2).

(2a) Peter was scratched by Greg at the car wash. As a result, he `got mad and made Greg pay for the car wash.` [referent of pronoun = subject]

(2b) Kate was punched by Lisa at the restaurant. As a result, she `got a bruise on her arm …` [referent of pronoun = subject]

(2c) Lisa pushed Anne on the beach. As a result, she `hurt her leg and got angry at Lisa.` [referent of pronoun = object]

(2d) Kate hit Mary at school. As a result, she got `angry at Kate.` [referent of pronoun = object]

(2e) John punched Peter at the museum. As a result, he `accidentally broke a vase.` [referent of pronoun = unclear]

3.3 Results and discussion

As can be seen in Figure 3, participants mostly interpreted the pronoun prompt as referring to the character that is the patient in the prompt sentence, i.e., the object in active voice and the subject in passive voice. In the active conditions, the *object preference* is significantly higher than chance ($t1(23)=2.34$, $p<.05$, $t2(23)=2.98$, $p<.01$) and in the passive conditions, the *subject preference* is significantly higher than chance ($t1(23)=6.72$, $p<.001$, $t2(23)=9.28$, $p<.001$). (In these analyses, we are looking only at the binary choice between subject continuations and object continuations, and thus chance was estimated to be 0.5.)

Figure 2 also suggests that the bias to interpret the pronoun as referring to the preceding patient is stronger in passives than in actives. To test this, we computed a

'patient advantage score' for both passives and actives, by subtracting the proportion of non-patient continuations from the proportion of patient trials both by subjects and by items. Paired t-tests show that sentences in the active voice do indeed result in a weaker patient preference than sentences in the passive voice ($t1(23)=-2.03$, $p=.054$, $t2(23)=-2.97$, $p<.01$). We suggest that the 'patient boost' exhibited by passives may be due to patient being promoted to a non-canonical subject position.

Overall, we find a clear effect of thematic role on pronoun interpretation: When comprehenders encounter a gender-ambiguous pronoun, they tend to interpret it as referring to the preceding patient. This preference arises both with patients in object position (actives) and patients in subject position (passives). Interestingly, we do not find a main effect of syntactic position/subjecthood, contrary to Fukumura & van Gompel's [8] findings with Stimulus-Experiencer and Experiencer-Stimulus verbs.

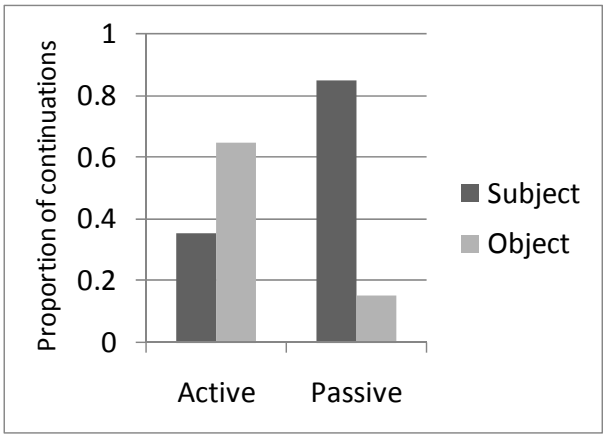

Fig. 3. How frequently did people use the prompt pronoun to refer to the preceding subject vs. the preceding object? (Experiment 1)

4 Experiment 2: No Prompt Pronoun

4.1 Participants, materials and design, coding

Experiment 2 had the same design as Experiment 1, except that no pronoun prompt was provided: the sound files were truncated at the end of the connective. The stimuli were otherwise identical to Experiment 1. Thus, participants were now free to continue the sentence in a more open-ended way and could use a range of referring expressions. Experiment 2 allows us to test (i) which referent participants are mostly likely to talk about, and thereby obtain a measure of predictability/likelihood-of-mention, and (ii) what kind of referring expressions participants choose when referring to different antecedents. Because participants had to produce the referring

expression themselves, we can also measure their durations to test for potential effects of acoustic reduction. Twenty-four new participants took part in Experiment 2.

Two coders worked through the data independently, working in the same way as described for Experiment 1. They noted what kind of referring expression participants used (e.g. pronoun, name, full noun phrase) and what it referred to (subject, object, unclear, another referent). Participants produced pronouns on about 11% of all trials,[2] and out of these, 27% were coded as having unclear/ambiguous antecedents. These unclear trials are excluded from subsequent analyses.

Ex.(3) shows examples of different continuations and how they were coded. In addition to the forms and coreferential relations in the continuations, the acoustic durations of the names produced by the participants were measured.

(3a) Kate hit Mary at school. As a result, `Mary did not want to be friends with Kate anymore.` [coded as form=name, referent=object]

(3b) Mike pushed Peter at the airport. As a result, `security took them both and talked to them.` [coded as form=noun, referent=other]

(3c) Lisa punched Kate at the restaurant. As a result, `they were banned from the restaurant for being overly dramatic.` [coded as form=plural pronoun, referent=subject & object]

(3d) Greg scratched Peter at the car wash. As a result, `he had marks on his arm for a week.` [coded as form=pronoun, referent=object]

(3e) Mary tapped Anne in the forest. As a result, `she woke up a bear and they had to run away.` [coded as form=pronoun, referent=unclear]

4.2 Results and discussion: Who is mentioned and how

We first report which entities are most likely to be mentioned and with what kind of referring expressions. Then, we turn to the acoustic durations of names produced by participants, to see whether predictability is associated with acoustic reduction.

Overall, participants produced a variety of continuation types.[3] When we look at only those trials where participants chose to start their continuation with a pronoun

[2] 11% is relatively low, and we often obtain higher rates of pronoun production in comparable written tasks where participants continue sentences without a visual scene, which suggests the low rate may be task-related. However, as we will see below, although the rate of pronouns in Exp.2 is low, the trials where people *did* produce pronouns show patterns that are strikingly similar to the interpretation of pronouns in Exp.1 (cf. Fig.3, Fig.4). This is a positive finding, because it shows that although the task may have led to an overall decrease in the proportion of pronouns produced, it crucially did not distort the underlying patterns.

[3] We focus on trials where the participants start by referring to the preceding subject or object. Continuations that start with other entities (e.g. *the restaurant manager*) or refer to both subject and object (e.g. *they*) are excluded, as are continuations where it was unclear whether people were referring to the subject or the object. This lead to 27% of trials being excluded.

(Figure 4), the patterns look very similar to Experiment 1 (Figure 3). In other words, we again find indications of an overall *patient preference*, which is stronger with passives. (Due to the low number of observations, no statistics are reported.) In other words, when participants chose to produce a pronoun, this mostly occurred when they were talking about the patient of the preceding sentence.

Let us now look more broadly at which referent participants are more likely to continue with, independent of the referring expression that is used. These results will give us a measure of which referents are most predictable, most likely to be mentioned next, in the different conditions. Here, a very different pattern emerges. As can be seen in Figure 5 (on the next page), the patient is *not* consistently the entity with the highest likelihood-of-mention:

In *active conditions*, most continuations do indeed start with the preceding object/patient, and the proportion of subject continuations is significantly lower than chance ($t1(23)=-6.33$, $p<.001$, $t2(23)=-8.82$, $p<.001$). Thus, after an active sentence followed by a causal connective, the patient has a high likelihood-of-mention.

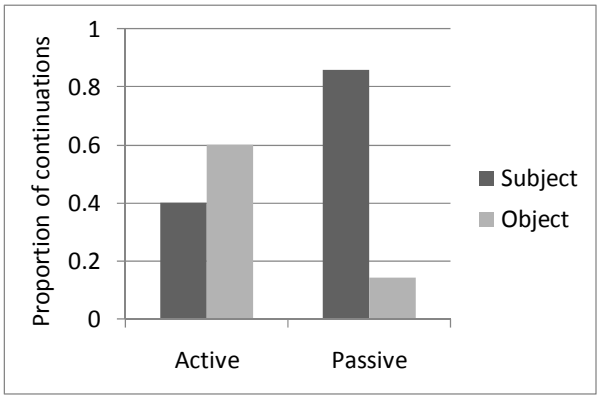

Fig. 4: On trials where participants chose to produce a pronoun, what did the pronoun refer to? (Experiment 2)

However, in the *passive conditions*, continuations were evenly split between the patient (the promoted subject) and the agent (demoted to the *by*-phrase), as can be seen in Figure 5 on the next page. The proportion of subject continuations does not differ significantly from chance ($t1(23)=.68$, $p>.5$, $t2(23)=.74$, $p>.4$). (In these analyses, as in Experiment 1, we only looked at those trials where people continued with either the subject or the object. I.e., we are focusing on the binary choice between subject vs. object, and thus chance was estimated to be 0.5.)

As a whole, we see that in passives, the patient and the agent are equally predictable, equally likely to be mentioned in subsequent discourse: Neither is a clear winner. Thus, the passive results reveal a striking dissociation between likelihood-of-mention and likelihood-of-pronominalization: Patients are the most likely to be pronominalized, but not the most likely to be mentioned next.

It is worth noting that the difference in the likelihood-of-mention patterns exhibited by actives and passives emerges from our results, i.e., it was not specifically manipulated as an independent variable. We regard this as a strength of our design, as

it means that the predictability differences exhibited by agents and patients depending on voice emerged naturally and were not pre-specified by the experimental design.

Similar to Experiment 1, we also computed patient advantage scores by subjects and by items. Paired t-tests show – not surprisingly, given the patterns visible in Figure 5 – that sentences in the active voice now have in a stronger patient preference than sentences in the passive voice ($t1(23)=3.65$, $p<.005$, $t2(23)=6.04$, $p<.0001$).

This pattern can also be described in terms of there being relatively more *agent* continuations in the passives than in the actives: Passives fail to show the overwhelming patient preference that we see in actives. In the next section, we consider the reasons for the boost in the rate of agent continuations with passives.

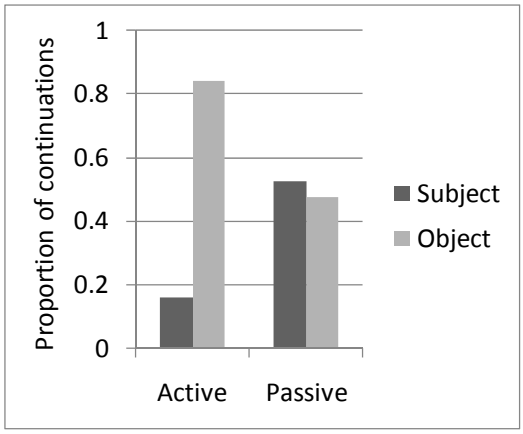

Fig. 5. How frequently did people start their continuations by reference to the preceding subject vs. the preceding object (regardless of form)? (Experiment 2)

4.3 Why do actives and passives differ with respect to likelihood-of-mention?

We saw in the preceding section that (i) active agent-patient verbs followed by a causal connective create an expectation that the patient will be mentioned next, but (ii) passive agent-patient verb render both the agent and the patient equally expected/predictable. Another way of putting this is to say that the passives boost the rate of agent continuations. An important question is why this difference arises.

One possibility is that the active/passive asymmetry can be derived from a patient-bias combined with linear recency effects: In active sentences, thematic role and linear recency coincide (the object is a patient, i.e., focused by causal connective, and also linearly the most recent argument), whereas in passive sentences, thematic role and recency favor different arguments (the grammatical subject and the object in the *by*-phrase, respectively). Could this be the reason why passives do not result in a clear expectation for one referent over the other?

We would like to suggest that the answer is no: We argue that the asymmetrical behavior of actives and passives cannot be reduced to linear recency. Evidence for this claim comes from other experiments, parallel to Experiments 1 and 2, that we conducted (see [16]) with stimulus-experiencer verbs (e.g. *annoy, irritate*). With

stimulus-experiencer verbs, a causal connective is known to focus attention on the experiencer (e.g. *Peter irritated John and as a result John stomped out of the room*, [12], [17]). Crucially, we found that with this verb class, actives and passives pattern in the same way: Participants preferred to start their continuations by talking about the experiencer, and this effect was not influenced by linear recency.

In light of these findings, instead of blaming the active/passive difference exhibited by agent-patient verbs on linear recency, we hypothesize that it stems from the thematic structure of agent-patient verbs. More specifically, agents are commonly regarded as semantically more prominent than patients (e.g. [18] and many others) and thus passivization of an agent-patient verb ends up demoting a semantically prominent argument (the agent) into a syntactically low-prominence position (the *by*-phrase of the passive). Thus, passivized agent-patient verbs can be regarded as suffering from a syntax-semantics mismatch. (Crucially, no such mismatch occurs when stimulus-experiencer verbs are passivized: the stimulus ends up in the *by*-phrase, but the thematic role of stimulus is commonly viewed as less prominent than the role of experiencer, and thus passivized stimulus-experiencer verbs do not suffer from a syntax-semantics mismatch).

In sum, when it comes to agent-patient verbs, we suggest that the asymmetrical behavior of actives and passives stems from the syntax-semantics mismatch that arises when semantically-prominent arguments – agents – are placed in a syntactically low-prominence position, the *by*-phrase.

However, it is worth noting that our findings regarding the dissociation between likelihood-of-mention and likelihood-of-pronominalization do not hinge on this question of *why* actives and passives differ. While the question of why they differ in this way is an important question in its own right, what is crucial for our claims regarding the relationship between likelihood-of-mention and likelihood-of-prominalization is simply that actives and passives do differ in terms of which argument is most likely to be mentioned subsequently, and that these likelihood patterns do not match the pronominalization patterns.

4.4 Results and discussion: Acoustic duration of names

Given the dissociation between likelihood-of-mention and likelihood-of-pronominalization, does this mean that predictability plays no role in reference-tracking? Or might we find predictability effects in another dimension?

In an initial exploration of this question, we looked at the duration of names that participants produced in subject position in Experiment 2 (e.g. the second occurrence of 'Lisa' (underlined) in *Mary slapped Lisa at the zoo. As a result Lisa stormed off angrily*). We compared names produced in the active condition, where the object of the prompt-sentence has a high likelihood of subsequent mention (i.e. it's the expected referent), and names produced in the passive condition, where there is no clearly expected referent, since continuations were split between subject and object (Fig. 4). All of the analyzed names were in subject position, following the connective 'As a result.' The key difference is simply whether the preceding sentence was active or passive. This allows us to investigate whether predictability influences duration of proper names. The durations of the names (from the onset of each name to the offset

of each name) were measured by a phonetically-trained annotator using Praat (Boersma & Weenink, *praat.org*). Datapoints whose duration was more than 4 standard deviations from the mean were excluded from analysis (1.6% of the data).

Analyses of duration show that names in subject position are shorter (283ms) after active sentences than names in subject position (320ms) after passive sentences (Figure 6). This finding is marginal by subjects and significant by items ($t1(22)=-1.8$, $p=.0858$, $t2(23)=-2.335$, $p<.03$).[4]

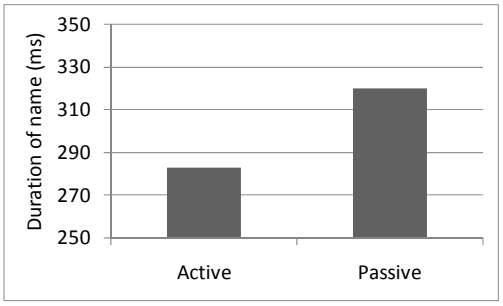

Fig. 6. Duration of names in subject position after actives and passives

Recall that in all conditions, the relevant name had been mentioned in the immediately preceding sentence. Thus, this effect cannot be attributed to whether the name was given or new information: all names were discourse-old/given. Additionally, all names occurred in both conditions and in each condition as both the agent and patient in different trials, and thus these results are not easily attributable to inherent differences in the names themselves.

In sum, we found that when participants produced names, the duration of the names showed effects of predictability. Although further work and more analyses are needed to assess the validity of our preliminary findings, we would like to suggest that effects of likelihood-of-mention/predictability do exist in the domain of reference-tracking, although they do not appear to influence referring-expression choice. An intriguing question for future work concerns the prosodic realization of the names, in particular the presence/absence of pitch accents related to contrast or newness. We plan to conduct further analyses to explore this issue, as it could provide additional evidence for effects of predictability in the domain of reference-tracking.

5 Conclusions

Findings from a range of domains suggest that predictability is an important aspect of human language processing. Our findings suggest that predictability can also influence referential processing, in the form of reduced acoustic durations for predictable referents. However, we do not find evidence that predictability is directly linked to the likelihood of pronominalization, contrary to what some researchers have claimed. Instead, our findings reveal that – at least in the contexts we tested – the use

[4] The by-subjects analysis has 23 participants, because one person did not start with any names.

and interpretation of pronouns is influenced by thematic role (e.g. agent vs. patient), independently of which referent is most predictable (i.e., most likely to be mentioned next). On the basis of our divergent findings for actives and passives, we also suggest that referents' likelihood-of-mention is influenced by an interplay of syntactic and semantic factors, in particular the mapping between syntactic roles and thematic roles. As a whole, our results highlight the benefits of exploring not only lexical but also acoustic aspects of referential production.

Acknowledgments. This work was partially funded by NIH grant 1R01HD061457 to E. Kaiser. Many thanks to Jackie Kim, Monica Do and Heidi Mettler for assistance with data analysis and coding. We also thank three anonymous reviewers for their useful comments and feedback. An earlier version of this project was presented at the *Workshop on the Production of Referring Expressions* (PRE-CogSci 2011).

References

1. Bell, A. et al.: Effects of disfluencies, predictability, and utterance position on word form variation in English conversation. *Journal of the Acoustic Society of America*, 113, pp. 1001--1024 (2003)
2. Jurafsky, D., Bell, A., Gregory, M., Raymond, W.: Probabilistic relations between words. *Frequency and the emergence of linguistic structure* (pp. 229--254). John Benjamins. (2001)
3. Hale, J.: A Probabilistic early parser as a Psycholinguistic Model. In *Proceedings of NACCL, 2,* pp. 159--166. (2001)
4. Levy, R.: Expectation-based syntactic comprehension. *Cognition*, 106, 1126--1177. (2008)
5. Hume, E., Mailhot, F.: The role of entropy and surprisal in phonologization and language change. In A. Yu, (ed.), *Origins of Sound Patterns*. University Press. (To appear)
6. Arnold, J.: Reference production: Production-internal and addressee-oriented processes. *Language and Cognitive Processes*, 23, pp. 495--527. (2008)
7. Givón, T.: *Mind, Code and Context: Essays in Pragmatics.* Lawrence Erlbaum. (1989).
8. Fukumura, K., van Gompel, R.: Choosing anaphoric expressions: Do people take into account likelihood of reference? *Journal of Memory and Language*, 62, pp. 52--66. (2010)
9. Kehler, A., Kertz, L., Rohde, H., Elman, J.: Coherence and coreference revisited. *Journal of Semantics*, 25, pp. 1--44. (2008)
10. Arnold, J.: The Effect of Thematic Roles on Pronoun Use and Frequency of Reference Continuation. *Discourse Processes*, 31(2), pp. 137--162 (2001)
11. Kaiser, E.: Investigating the Consequences of Focus on the Production and Comprehension of Referring Expressions. *International Review of Pragmatics*, 2, pp. 266--297. (2010)
12. Rohde, H.: *Coherence-Driven Effects in Sentence and Discourse Processing*. Ph.D., UC San Diego. (2008)
13. Kehler, A.: *Coherence, Reference, and the Theory of Grammar*. CSLI Publications. (2002)
14. Asher, N., Lascarides, A.: *Logics of Conversation*. Cambridge University Press. (2003)
15. Kertz, L., Kehler, A., Elman, J.: Evaluating a Coherence-Based Model of Pronoun Interpretation, *Workshop on Ambiguity in Anaphora,* Malaga, Spain. (2006)
16. Kaiser, E., Holsinger, E., Li, D.: *Who next? Consequences of syntax-semantics mismatches for likelihood of mention.* Poster, CUNY Conf on Human Sentence Proc., Stanford. (2011)
17. Stevenson, R., Crawley R., Kleinman D.: Thematic roles, focusing and the representation of events. *Language and Cognitive Processes* 9, pp. 519--548. (1994)
18. Grimshaw, J.: *Argument Structure*. Cambridge: MIT Press. (1990)