Journal of Memory and Language 87 (2016) 128–143

Contents lists available at ScienceDirect

Journal of Memory and Language

journal homepage: www.elsevier.com/locate/jml

Talker-specificity and adaptation in quantifier interpretation

Ilker Yildirim [a,b,c,*], Judith Degen [d], Michael K. Tanenhaus [a,e], T. Florian Jaeger [a,e,f,*]

[a] University of Rochester, Department of Brain and Cognitive Sciences, United States
[b] Massachusetts Institute of Technology, Department of Brain and Cognitive Sciences, United States
[c] The Rockefeller University, Laboratory of Neural Systems, United States
[d] Stanford University, Department of Psychology, United States
[e] University of Rochester, Department of Linguistics, United States
[f] University of Rochester, Department of Computer Science, United States

ARTICLE INFO

Article history:
Received 6 August 2014
revision received 7 August 2015
Available online 11 November 2015

Keywords:
Adaptation
Talker-specificity
Quantifiers
Semantics
Pragmatics

ABSTRACT

Linguistic meaning has long been recognized to be highly context-dependent. Quantifiers like *many* and *some* provide a particularly clear example of context-dependence. For example, the interpretation of quantifiers requires listeners to determine the relevant domain and scale. We focus on another type of context-dependence that quantifiers share with other lexical items: talker variability. Different talkers might use quantifiers with different interpretations in mind. We used a web-based crowdsourcing paradigm to study participants' expectations about the use of *many* and *some* based on recent exposure. We first established that the mapping of *some* and *many* onto quantities (candies in a bowl) is variable both within and between participants. We then examined whether and how listeners' expectations about quantifier use adapts with exposure to talkers who use quantifiers in different ways. The results demonstrate that listeners can adapt to talker-specific biases in both how often and with what intended meaning *many* and *some* are used.

Introduction

The meaning of many, if not all, words is context-dependent. For example, whether we want to say that John is *tall* depends on whether John is being compared to other boys his age, professional basketball players, dwarves, etc. (e.g., Halff, Ortony, & Anderson, 1976; Kamp, 1995; Kennedy & McNally, 2005; Klein, 1980). Other words whose interpretation requires reference to context are pronouns and quantifiers (Bach, 2012). For example, the interpretation of a quantifier like *many* depends on the class of objects that is being quantified over: the number of crumbs that *many crumbs* refers to is judged to be higher than the number of mountains that *many mountains* refers to (Hörmann, 1983).

A less-well studied aspect of context-dependence is how a given talker uses quantifiers like *many* and *some*. Talkers exhibit individual variability at just about any linguistic level investigated – including, for example, pronunciation (e.g., Allen, Miller, & DeSteno, 2003; Bauer, 1985; Harrington, Palethorpe, & Watson, 2000; Yaeger-Dror, 1994), lexical preferences (e.g., Finegan & Biber, 2001; Roland, Dick, & Elman, 2007; Tagliamonte & Smith, 2005), and syntactic preferences (e.g., the frequency with which they use passives, Weiner & Labov, 1983). Therefore, talkers are also likely to differ in how they use quantifiers. For example, talkers may differ in how many crumbs they consider to be *many crumbs*, and these differences would consequently be reflected in their productions. In this case, listeners would be well served by taking into account

* Corresponding authors at: University of Rochester, Department of Brain and Cognitive Sciences, United States.
E-mail addresses: ilkery@mit.edu (I. Yildirim), fjaeger@bcs.rochester.edu (T.F. Jaeger).

http://dx.doi.org/10.1016/j.jml.2015.08.003
0749-596X/

talker-specific knowledge in order to successfully infer what the talker intended to convey.

Talker-specific knowledge has been observed experimentally in cases of variation in pronunciation and syntactic production (e.g., Clayards, Tanenhaus, Aslin, & Jacobs, 2008; Creel & Bregman, 2011; Creel & Tumlin, 2009; Fine, Jaeger, Farmer, & Qian, 2013; Kamide, 2012; Kraljic & Samuel, 2007). While this question has received less attention in lexical processing, there is some evidence that listeners can learn to anticipate talker-specific biases in the frequency with which referents are being referred to (Metzing & Brennan, 2003) and that these talker-specific expectations are reflected in online processing (e.g., Creel, Aslin, & Tanenhaus, 2008). These studies complement classic work on *conceptual pacts* in which interlocutors adjust their use of referential expressions to create temporary, shared context-specific names (Brennan & Clark, 1996).

Previous work on talker-specific lexical expectations has focused on open class, semantically rich, content words – typically nouns (Brennan & Clark, 1996; Creel et al., 2008; Metzing & Brennan, 2003). This raises the question of whether listeners are capable of adapting to talker-specific differences in the use of words that convey more abstract meanings, such as those of quantifiers. If listeners do in fact adapt to talker-specific differences, what specifically are listeners adapting to, i.e., what is the nature of the representations that are being updated and what are the underlying mechanisms?

The current paper begins to address these questions by studying adaptation to talker-specific differences in the use of the quantifiers *some* and *many*. We present four experiments that investigate lexical adaptation. Taken together, these experiments establish (i) that listeners *can* adapt to talker-specific differences in the usage of even abstract lexical items, such as quantifiers; (ii) that, provided sufficient exposure, such adaptation can be achieved even for multiple talkers simultaneously; (iii) that lexical adaptation is observed both to talker-specific differences in the frequency with which lexical items are used *and* to talker-specific differences in how they are being used; and thus, finally, (iv) that lexical adaptation – although often studied as a separate phenomena—exhibits many of the hallmarks of adaptation observed for other linguistic domains. Next, we elaborate on these points, while introducing the four experiments presented below. In doing so, we relate our research to previous work and highlight the contributions of the current work.

Before we investigate lexical adaptation to talker-specific quantifier use, we first assess whether the premise for adaptation is given: Experiment 1 demonstrates that listeners differ in their initial expectations about a talker's use of a variety of quantifiers, including *some* and *many*. This shows that if listeners want to arrive at an interpretation of an utterance that is close to the talker's intended meaning, they might sometimes need to adapt their expectations about quantifier use to match those of the current talker. Experiment 1 thus provides the first direct evidence that there would potentially be a *benefit* to adaptation to talker-specific differences in quantifier use.

This then raises the question whether listeners do adapt to these changes. This is the central motivation for Experiment 2. Going beyond this question and previous work, Experiment 2 also begins to investigate the nature of the changes in expectations that result from exposure to a novel talker. Specifically, we ask whether lexical adaptation can be talker-specific. The answer to this question is of theoretical relevance, as it speaks to the nature of the mechanisms underlying lexical adaptation. We briefly elaborate on this point, as it has so far received relatively little attention in the literature on lexical adaptation (but see Brennan & Clark, 1996; Pickering & Garrod, 2004).

A priori, there are several ways in which a listener can treat experience with a novel talker. A listener might treat new experience as evidence that can be used to sharpen prior expectations about quantifier use without taking into account the specific context, including the talker. Any adaptation would then be to talkers in general. At the other extreme, adaptation might be completely context-specific. If that were the case, then adaptation would be specific to a particular talker in a particular context and would not at all generalize to other talkers. A more likely possibility is that listeners strike a subtle balance between context-general and context-specific adaptation (cf. Kleinschmidt & Jaeger, 2015). Prima facie, it would seem undesirable for a language processing system to allow a small amount of recent exposure to overwrite life-long experience with language. At the same time, it is beneficial to be able to rapidly adapt to talker-specific lexical preferences, potentially increasing the efficiency of communication (for related discussion, see Brennan & Clark, 1996; McCloskey & Cohen, 1989; McRae & Hetherington, 1993; Pickering & Garrod, 2004; Seidenberg, 1994).

One way to meet both the need for adaptation and the need to maintain previously acquired knowledge is to learn and maintain talker-*specific* expectations, so that adaptation to a novel talker does not imply loss of previously acquired knowledge. Research in speech perception has explored and found support for this hypothesis (Goldinger, 1996; Johnson, 2006; Kraljic & Samuel, 2007; for review, see Kleinschmidt & Jaeger, 2015). More recent research has found support for this idea in other domains of language processing (e.g., prosodic processing, Kurumada, Brown, Bibyk, Pontillo, & Tanenhaus, 2014; Kurumada, Brown, & Tanenhaus, 2012; and sentence processing, Fine et al., 2013; Jaeger & Snider, 2013). For example, in episodic and exemplar-based models, linguistics experiences are assumed to be stored along with knowledge about the context in which they occurred (Goldinger, 1996; Johnson, 2006; Pierrehumbert, 2001). This is how these models capture talker-specific expectations. (Similar reasoning applies to Bayesian models of adaptation that assume generative processes over hierarchically organized indexical alignment, Kleinschmidt & Jaeger, 2015). Similarly, memory-based models of lexical alignment (Horton & Gerrig, 2005, in press) can in theory account for both talker-specific expectations – if talkers are included as contexts (Brown-Schmidt, Yoon, & Ryskin, 2015).

Changes in the use of lexical forms and structures due to exposure are often attributed to temporary changes in expectation within a spreading-activation framework. These "priming-based" accounts assume that exposure