Predictions for self-priming from incremental updating models unifying comprehension and production

Cassandra L. Jacobs
Department of Psychology
University of Illinois at Urbana-Champaign
603 E. Daniel St.
Champaign, IL 61820
cljacob2@illinois.edu

Abstract

Syntactic priming from comprehension to production has been shown to be robust: we are more likely to repeat structures that we have previously heard. Many current models do not distinguish between comprehension and production. Here we contrast human language processing with two variants of a Bayesian belief updating model. In the first model, production-to-production priming (i.e. self-priming) is as strong as comprehension-to-production priming. In the second, both individuals who self-prime and those who do not are exposed to a syntactic construction via comprehension. Our results suggest that when production-to-production priming is as robust as comprehension-to-production priming, then speakers who self-prime are simultaneously less likely to be primed by input from comprehension and demonstrate different distributions of responses than speakers who do not self-prime. The computational model accords with recent results demonstrating no self-priming, and provides evidence for an account of syntactic priming that distinguishes between production and comprehension input.

1 Introduction

Syntactic priming in production is the increased probability of using a syntactic structure when we have recently encountered it. Many models of syntactic priming treat comprehension and production as equally influential on the language production system (Pickering & Garrod, 2013; Chang et al., 2006; Reitter et al., 2011). This means that syntactic priming can occur without another person being present, from production to production (self-priming). Unfortunately, little experimental research has assessed the degree to which priming occurs from production to production in a controlled context.

The primary phenomena that are of interest to research on syntactic priming are as follows: (1) Do we incrementally and cumulatively adapt to our linguistic environment, changing how we talk based on what we hear (Jaeger & Snider, 2013; Chang et al., 2006; Kaschak et al., 2012; Reitter et al., 2011; Pickering & Garrod, 2013)? (2) Do we change how we talk more when we encounter less probable structures (Jaeger & Snider, 2013; Chang et al., 2006)? (3) How long-lasting is syntactic priming (Kaschak, 2007; Bock, 1986)? And finally: (4) Are we more likely to repeat the structures that we ourselves have said, as predicted by models that unify priming in comprehension and production (Pickering & Garrod, 2013; Chang et al., 2006)? The model presented here accounts for these phenomena and makes additional predictions about the potential ramifications of self-priming on linguistic representations and efficient communication.

2 Psycholinguistic evidence

In any model where speakers are influenced by their prior experience, regardless of the source, priming should occur. Corpus linguistics has provided evidence of self-priming in production. Some of these studies assessed priming of specific constructions such as the dative alternation or relative clauses

Proceedings of CMCL 2015, pages 1–8,
Denver, Colorado, June 4, 2015.

(Jaeger & Snider, 2013; Myslin & Levy, submitted) or instead model the whole syntactic system using probabilistic rules in several grammatical formalisms (Gries, 2005; Healey et al., 2014; Reitter et al., 2011). These models generally find evidence for syntactic priming from production to production, or structural repetition that occurs more often than would be expected by chance. In contrast to some of these results, some studies report that structural repetition occurs less than would be expected by chance (Healey et al., 2014). Healey et al. (2014) argued that this is because speakers avoid continuously reusing syntactic structures.

The evidence for self-priming in a more controlled environment is weaker. Within experimental psycholinguistics, structural priming in production is almost always mediated by some intervening comprehension task like sentence completion or a memory task (e.g. Kaschak, 2007, Bock, 1986, etc.), making it difficult or impossible to assess whether self-priming occurs. Some studies have reported that some speakers have an almost uniform bias toward one structure or another (e.g. Jaeger & Snider, 2013). However, despite these claims that comprehension and production show similar amounts of syntactic priming (Tooley & Bock, 2014), production has also been shown to be substantially less flexible than comprehension (Remez, 2013). Syntactic persistence within a speaker could simply be a product of an individual's own syntactic preferences rather than self-priming per se.

Some experimental work has been conducted to test whether comprehension and production are weighted equally in structural priming in production. Counter to what has been found in corpus studies, Jacobs et al. (2015) failed to find evidence for self-priming despite strong comprehension-to-production priming. In their study, participants produced 7 dative descriptions, comprehended 6 dative descriptions of a single form of the construction, and produced an additional 7 descriptions in order to identify effects of self-priming, individual differences, and comprehension input on the magnitude of comprehension-to-production priming. They found that the rates of structural repetition were flat across the experiment, but speakers were strongly sensitive to comprehension input, showing large and sustained comprehension-to-production priming effects. They also found larger priming effects for the less probable syntactic structure, consistent with error-driven learning accounts of the inverse frequency effect (Jaeger & Snider, 2013).

While the evidence for self-priming has been weaker, almost all syntactic priming studies have demonstrated that comprehension plays a very large role in production preferences. This is to be expected if priming is a means of achieving efficient communication, which requires using comprehended language to modify our own productions to be more easily understood (Pickering & Garrod, 2013; Tooley & Bock, 2014). It is less clear what the functional role of self-priming would be, however. If speakers are likely to repeat structures they have recently used, the language community may end up with two types of speakers: those who only use one structure or the other, which would pose some difficulty for comprehenders (MacDonald, 2013).

We aim to address the question of the equivalence of learning from comprehension and production, as well as account for the results of Jacobs et al. (2015). To do this, we constructed a very simple Bayesian belief-updating model that makes predictions about syntactic preferences with and without self-priming. This model is very similar to those of Fine et al. (2010), Kleinschmidt et al. (2012), and Myslin & Levy (submitted), who have modeled updating in syntactic comprehension. These models perform Bayesian belief-updating of the probabilities of outcomes in a syntactic alternation. In ours, we focus on the prepositional versus double object dative, though the model can be extended to any syntactic alternation. In the computational model, we demonstrate differences at the individual and population levels differences between self-priming and no self-priming (Model 1), as well as characterize what self-priming does to syntactic adaptation in production after comprehension (Model 2).

3 Model structure

Many of the recent computational models that have sought to account for syntactic priming effects in comprehension and production use incremental algorithms to represent trial-by-trial effects, while also treating experience in comprehension and production as contributing equally to the implicit learning process (Jaeger & Snider, 2013; Chang et al., 2006; Reitter et al., 2011; Pickering & Garrod, 2013). The model here is simple but relies on similar assumptions: representations are changed when language is processed.

Each utterance choice can be conceptualized as a single coin flipped randomly (i.e. a binomial process). We generate sequences of utterances by sampling from either a static Binomial distribution or one that is continually being updated. The structural alternation we consider here is the double object dative construction, which has been extensively studied in syntactic priming experiments (Jaeger & Snider, 2013; Kaschak, 2007; Bock, 1986; Pickering & Garrod, 2013). In this syntactic alternation, direct and indirect objects can change places after some English verbs like *give*, *hand*, *throw*, or *show*:

- The librarian gave the book to the boy. (prepositional object - PO)
- The librarian gave the boy the book. (double object - DO)

For every sentence a speaker intends to use a dative structure, the production system will select either the PO or the DO form of the sentence. Selections are single samples from a Binomial distribution with p(PO) = .55, which was derived from the empirical probability of individuals producing a PO in Jacobs et al. (2015).

Previous models have used error-driven learning to account for incremental adjustments in syntactic preferences (Jaeger & Snider, 2013; Chang et al., 2006). In our model, as in similar models, we update the probability of using a PO using the Beta prior, which is conjugate to the Binomial distribution. This prior is also appropriate because in any syntactic alternation, there are two possible outcomes. The Beta distribution has only two hyperparameters, α and β. α roughly corresponds to the number of times the model believes that it has experienced a PO dative, and β represents the analogous number for the DO dative.

The selection of the α and β values was based on exploring an integer-valued parameter space between 1 and 10 for both parameters. In Figure 1 we plot the amount of self-priming all 100 models demonstrated as a function of the two hyper-

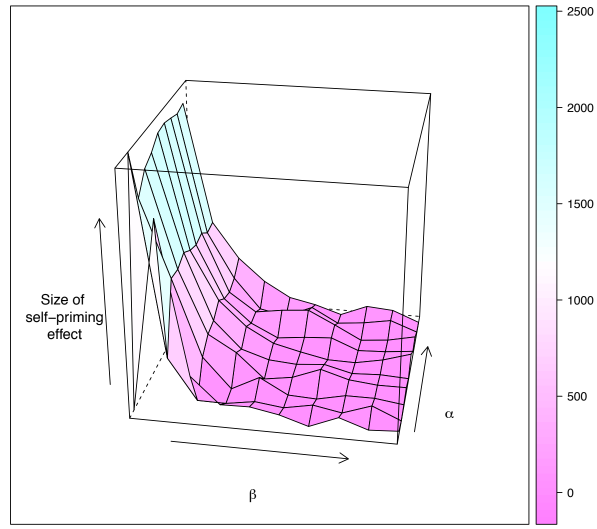

Figure 1: As values of α and β increase from 1 to 10, the models become less capable of self-priming. At the smallest values of β, the model is very sensitive to its own prior productions, while α does not seem to play a large role in self-priming.

parameters. We determined that α does not influence the sensitivity the model has to its own productions. Contrarily, β values on the smaller end tended to lead to large self-priming effects. Consequently we chose identical values for α and β that would allow for both self-priming and sensitivity to comprehension input. This was therefore a compromise between flexibility and rigidity.

We initialized the model with relatively weak, but not completely unbiased, priors that both structures were equally possible: $\alpha = 4$, $\beta = 4$. We assume that one of the structures is slightly more probable, with p(PO) = .55, as observed in Jacobs et al. (2015).

With every subsequent dative description that the model processes (either via production or comprehension), with k POs and n descriptions containing either a DO or a PO, the priors are adjusted via the following update rules:

$$\alpha = \alpha + k \tag{1}$$

$$\beta = \beta + n - k \tag{2}$$

The new probability of a PO is therefore:

$$p(PO) = \frac{\alpha + k + 1}{\alpha + \beta + n - 2} \tag{3}$$

For both Model 1 and Model 2, we ran 1000 “experiments” of 200 “participants” each. Model 1 looks at the production of seven double object datives to assess self-priming. Model 2 is structured like the experiments of Jacobs et al. (2015) and consists of three stages. First, seven datives are produced, then six datives of a single syntactic structure are comprehended, and then seven additional datives are produced. In the production tasks, participants' syntactic choices were randomly sampled from a Binomial distribution initialized at p(PO) = .55. This allows us to treat individual subjects differently prior to exposing them to the comprehension materials.

The model always updates the probability of a PO when it encounters a PO or DO. In Model 1 we assess what should happen to participants' later productions if they are influenced by their own productions (self-prime) to the same extent as in comprehension. In Model 2 we look at the magnitude of comprehension-to-production priming when participants are allowed to self-prime versus not.

4 Model 1 - Self-priming in production

This task can be conceptualized as a spontaneous production task. We conducted 1000 experiments with 200 participants each. Each participant produces seven sentences. Participants' syntactic choices are sampled from an incrementally updating Binomial distribution where the calculated posterior probability for each subject replaces the prior probability of p(PO). Each utterance that the model produces contains either a prepositional object dative construction (e.g. *The librarian gave the book to the boy*) or a double object dative (e.g. *The librarian gave the boy the book*).

Qualitatively, the model makes the prediction that, in general, models where self-priming occurs should be more likely to prefer one form of a structure over the other. Self-priming increases the entropy of the distribution, though it is naturally biased toward producing more probable structures since those are initially more likely to be drawn. We summarize the results of our simulations for self-primers and non self-primers below in Figure 2.

Importantly, self-priming always produces some change to the model parameters. At the limit, the production system's predictions are often correct,

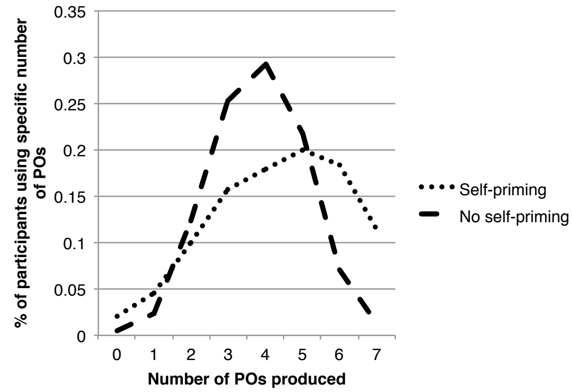

Figure 2: Self-priming without adjustment from comprehension results in a more uniform distribution of syntactic preferences across a population of speakers than in a model where speakers do not prime themselves.

leading to minimal prediction error (e.g. Chang, Dell, & Bock, 2006; Jaeger & Snider, 2013). A lack of prediction error can be conceptualized as the same thing as no self-priming in production, since neither leads to changes in syntactic preferences. Analogously, as the hyperparameters increase for an individual iteration of the model, the more confident the model is in its initial representations, leading to a stable distribution for a single speaker. It is also possible, however, that production and comprehension are updated separately, but production possesses more conservative hyperparameters, in line with research showing the static preferences of the production system (Remez, 2013).

The bulk of priming occurs early in the experiment. Participants' prior productions should play an influential role on their later productions, with the population's structural preferences stabilizing over the course of seven trials. This implies that prediction error made by the model decreases over many productions by a single speaker. We visualize this below in Figure 3.

5 Model 2 - Transfer from comprehension to production

We wanted to see how much self-priming diminishes the effect of comprehension-to-production priming, since self-priming changes the hyperparameters of the model, meaning that the model becomes more conservative in its estimates of the probability of us-

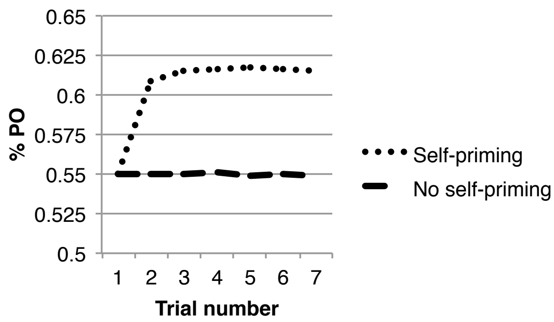

Figure 3: Models where self-priming is permitted converge on preferring the higher-probability structure within two trials. Models that do not permit self-priming show the same syntactic bias over the course of the experiment.

ing a PO (that is, the values of both α and β are larger). This experiment is structured in the same way as Model 1, except participants later comprehend 1 or 6 utterances containing a single structure (DOs or POs only) and produce an additional seven utterances in a final production task, following the experimental design of Jacobs et al. (2015).

Priming from comprehension to production is accomplished via the same updating process that we defined above. When a participant encounters a DO or PO structure, this updates the α and β parameters, and the observed posterior probability replaces the prior probability of a PO. The α and β parameters after comprehension are therefore adjusted in favor of the primed structure, which makes the primed structure more likely to be drawn on the next trial.

The model changes more for low-probability structures. Because the model has flexible priors, after hearing a single structure for six consecutive trials, the model will believe that the primed structure, regardless of which structure was primed, will be a certain likelihood (e.g. p(PO) or p(DO) = .75) after the comprehension stage. The model therefore shows more priming after exposure to the less common structure (DO) than after exposure to the more probable structure (PO). The model here demonstrates greater priming for dispreferred structures because the gap to bridge is larger for the dispreferred structure. Our results accord with models that employ error-based learning principles (Chang et al., 2006; Jaeger & Snider, 2013), where model parameters are adjusted via experience and change more in response to lower-probability structures. Figure 4 summarizes the relationship between prime probability, self-priming, and cumulativity.

Self-priming leads to much smaller effects of priming from comprehension. Even when exposed to a relatively high number of primes (6 versus 1), participants who self-prime do not align to comprehension input as much as participants who are not affected by their own prior productions. This is because the hyperparameters are tuned an additional seven trials before these participants go into the comprehension stage of the experiment, which makes it more difficult to change the probability distribution on later trials. Should the non self-priming models' hyperparameters be set sufficiently high, to perhaps the same level as the average self-priming model, this particular difference would likely disappear. Additionally, if hyperparameters for both models were even higher, it is likely that the magnitude of priming would not differ much between the two model types. Keeping the hyperparameters small before relevant linguistic experience (via comprehension or via both comprehension and production), makes it easier to see that self-primers require more comprehension input to offset their learned production biases.

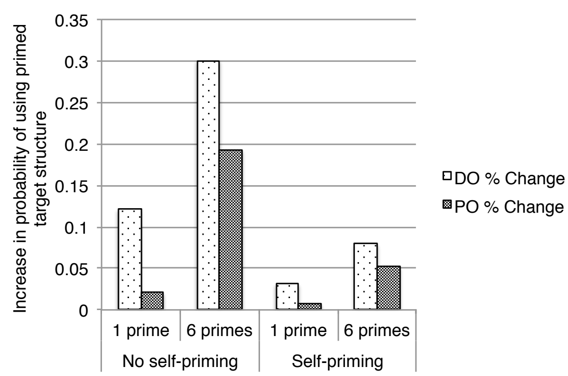

Figure 4: Increase in use of primed target structures in the second half of the experiment as a function of the number of primes, whether the models self-prime or not, and which structure was primed. All models accommodate the comprehension input. Priming is strongest when the primed structure was less frequent (DO). Similarly, when participants are exposed to more primes (6 versus 1), the priming effects are larger. Self-primers show smaller priming effects in all cases.

To assess the possibility that having high values of α and β diminishes the effect of comprehension-to-production priming, we ran an additional set of experiments with α and β set to 40 and 50, respectively. These values were selected to produce a probability of *p*(PO) roughly equal to .55. This experiment definitively demonstrated that an additional 6 primes of a particular structure still results in change to the model, though the non self-priming model is much more sensitive to six items of a single structure than the self-priming model, which has stronger beliefs about *p*(PO). The self-priming model seems inclined to not accommodate at all (.4% change), while the non self-priming model changes at a measurable amount (3%, comparable to the self-priming model that received 1 prime in Figure 4). The results of these simulations are demonstrated below in Figure 5. This suggests that even when the model's hyperparameters are made to be very conservative, self-priming is detrimental to comprehension-to-production priming.

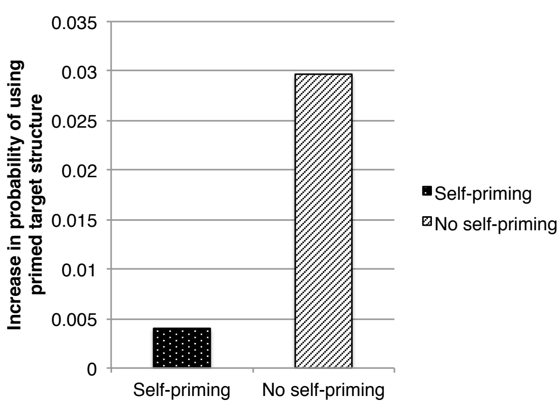

Figure 5: In comprehension-to-production priming models where self-priming is allowed to occur, the ability to accommodate novel comprehension input is greatly reduced. The self-priming model changes toward the comprehension input .4%, while the non self-priming model changes 3%.

6 Comparison to psycholinguistic data

An important test of the model is to compare its output to existing psycholinguistic data. While the model can account neither for the corpus linguistic studies that have demonstrated repetition greater than would be expected by chance (Jaeger & Snider, 2013; Reitter et al., 2011; Gries, 2005; Myslin & Levy, submitted), nor for those studies reporting speakers with strong structural biases (Jaeger & Snider, 2013), nor for anti-repetition and decay (Healey et al., 2014; Reitter et al., 2011), it can account for several phenomena that have been well-demonstrated in structural priming in production, as well as the results of Jacobs et al. (2015).

Self-priming is predicted by many accounts of syntactic priming where comprehension and production input are equated (Chang et al., 2006; Pickering & Garrod, 2013; Jaeger & Snider, 2013; Tooley & Bock, 2014). We have demonstrated here that self-priming leads to self-convergence. In a model where one's own prior productions weigh in on one's later productions, speakers quickly converge on a syntactic preference, with the majority of speakers coming to prefer the more probable structure (Figure 2). These results are counter to the experimental results of Jacobs et al. (2015), who reported that repetition of syntactic structures was constant at all stages of the production task. Had speakers primed themselves, the probability of using the more probable structure should have increased, however modestly (Figure 3).

Models allowing for self-priming with sufficiently large hyperparameters (Figure 1) would likely show both very little self-priming (having converged on a particular syntactic preference) as well as relatively little sensitivity to comprehension input (Figure 4). This is again counter to the results of Jacobs et al. (2015), who found very large comprehension-to-production structural priming effects. Production and comprehension sharing hyperparameters leads to radically different results than what is empirically observed.

Altogether, the computational model and the experimental data suggest that comprehension and production should be modeled separately. In order to obtain a production system that maintains its own internal representations *and* accommodates comprehension input (Remez, 2013; Pickering & Garrod, 2013), it is important that speakers not overweight their own syntactic preferences. This can be accomplished by having separate production and comprehension systems, where production-to-production priming (updating) is minimal while comprehension-to-production priming is large. An alternative, as described in Model 2, is to only allow updating from comprehension to production.

The model agrees with both existing computational models of syntactic priming and psycholinguistic data. The model updates incrementally, as many others have done (Chang et al., 2006; Jaeger & Snider, 2013; Fine et al., 2010; Kleinschmidt et al., 2012) and shows greater priming effects when exposed to multiple instances of the same structure (Kaschak et al., 2012). The priming effects are larger for less common structures than for more common structures because of the sensitivity of the model to error, demonstrating the inverse frequency effect (Chang et al., 2006; Jaeger & Snider, 2013). Finally, in a model where self-priming does not occur, or where updating from production to production is extremely low, structural preferences persist (Reitter et al., 2011; Tooley & Bock, 2014; Bock & Griffin, 2000; Kaschak, 2007; Kaschak et al., 2012).

7 Conclusion

The present model accounts for the incrementality, cumulativity, error sensitivity, and persistence of syntactic priming in production. In contrast to previous models of syntactic priming (Jaeger & Snider, 2013; Chang et al., 2006; Reitter et al., 2011; Pickering & Garrod, 2013), we tested the effects of equating comprehension and production input in structural priming in production. Self-priming has consequences for both individual and population-level language use.

This model makes predictions for all of these phenomena by making a single assumption: prior experience affects syntactic choices in production. Regardless of whether self-priming is allowed to occur or not, we are sensitive to recent and cumulative linguistic input and are primed to produce the structures we hear. Additionally, we change our representations more when we encounter low-probability structures. Because syntactic representations are updated without respect to time, syntactic priming effects do not necessarily decay in this model as in others (e.g. Reitter et al., 2011). However, most syntactic priming studies report structural persistence, making the model consistent with such studies (Bock, 1986; Tooley & Bock, 2014; Bock & Griffin, 2000; Kaschak et al., 2012). Finally, if individuals are allowed to self-prime, priming from comprehension will be weak and a population of speakers may be very variable in their structural preferences.

The functional value of self-priming is tempting if language is structured optimally to be easy for the speaker (MacDonald, 2013). Speakers can employ their own syntactic preferences, with comprehenders accommodating them. Other theories have stated that priming between speakers is necessary for efficient communication (Pickering & Garrod, 2013; Jaeger & Snider, 2013), especially because speakers must learn the distributions of the language around them in order to become successful communicators (Chang et al., 2006). In a language community, it may be sufficient to accommodate conversation partners rather than to develop highly idiosyncratic production preferences, making self-priming and structural repetition sub-optimal (Healey et al., 2014). At the same time, comprehenders are sensitive to the repetitive nature of conversations and may come to expect repetition during dialogue (Myslin & Levy, submitted). To that end, we are extending this model to simulate population dynamics in communities where speakers prime each other and possibly also prime themselves. Self-priming has consequences for language-level statistics, so it is important to see what changes might take place in a language community where individual speakers are allowed to become highly idiosyncratic.

It is still an open question as to whether one's own productions influence later syntactic choices. Some more recent psycholinguistic evidence suggests that they do not (Jacobs et al., 2015). If speakers do not self-prime, is it because they tend to avoid repeating syntactic structures (e.g. Healey et al., 2014), or is there an error-driven learning component, where predictions about one's own production are almost always correct, leading to no learning? To answer these questions, further experimental work is needed. In the meantime, we have outlined some predictions and accounted for the vast majority of phenomena in syntactic priming in production with a very simple belief-updating model.

References

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, *18*, 355–387.

Bock, K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General*, *129*, 177–192.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, *113*, 234–272.

Fine, A. B., Qian, T., Jaeger, T. F., & Jacobs, R. A. (2010). Syntactic adaptation in language comprehension. *Proceedings of ACL Workshop on Cognitive Modeling and Computational Linguistics*, (pp. 18–26).

Gries, S. T. (2005). Syntactic priming: A corpus-based approach. *Journal of Psycholinguistic Research*, *34*, 365–399.

Healey, P. G. T., Purver, M., & Howes, C. (2014). Divergence in dialogue. *PLoS ONE*, *9*, e98598.

Jacobs, C. L., Bock, K., & Watson, D. G. (2015). Speakers do not self-prime. *Poster presented at CUNY Conference on Sentence Processing, University of Southern California, Los Angeles.*.

Jaeger, T. F., & Snider, N. E. (2013). Alignment as a consequence of expectation adaptation: Syntactic priming is affected by the prime's prediction error given both prior and recent experience. *Cognition*, *127*, 57–83.

Kaschak, M. P. (2007). Long-term structural priming affects subsequent patterns of language production. *Memory & Cognition*, *35*, 925–937.

Kaschak, M. P., Kutta, T. J., & Coyle, J. M. (2012). Long and short term cumulative structural priming effects. *Language, Cognition and Neuroscience*, *29*, 728–743.

Kleinschmidt, D., Fine, A. B., & Jaeger, T. F. (2012). A belief-updating model of adaptation and cue combination in syntactic comprehension. *Proceedings of the 34rd Annual Meeting of the Cognitive Science Society (CogSci12)*, (pp. 605–610).

MacDonald, M. C. (2013). How language production shapes language form and comprehension. *Frontiers in Psychology*, *4*.

Myslin, M., & Levy, R. (submitted). Comprehension priming as rational expectation for repetition: Evidence from syntactic processing. *Cognition*.

Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, *36*, 329–347.

Reitter, D., Keller, F., & Moore, J. D. (2011). A computational cognitive model of syntactic priming. *Cognitive Science*, *35*, 587–637.

Remez, R. E. (2013). Analogy and disanalogy in production and perception of speech. *Language, Cognition and Neuroscience*, *30*, 273–286.

Tooley, K. M., & Bock, K. (2014). On the parity of structural persistence in language production and comprehension. *Cognition*, *132*, 101–136.