

Informativity and accessibility in incremental production of the dative alternation

Neil Rath¹, Brandon Waldon², and Judith Degen³

¹Stanford University, rathi@stanford.edu

²Georgetown University, bw686@georgetown.edu

³Stanford University, jdegen@stanford.edu

Abstract

Variation in the use of syntactic alternations has long been an explanatory target of language production theories. In this work, we test the predictions of several semantic, pragmatic and psycholinguistic theories of language use for the English dative alternation. We first experimentally test the role of incremental constituent informativity in the dative alternation, and find that contrary to information structural and RSA models of production, informativity has little effect on production preferences. We then more rigorously focus on accessibility effects, demonstrating that a lossy-context automatic policy can recover a key pattern of accessibility. Ultimately, we conclude that audience design pressures likely do not influence incremental production, but simply may affect planning at a broader scope.

Keywords: informativity; production; experimental methods

Introduction

A key challenge for theories of language production is to explain how speakers choose among multiple roughly meaning-equivalent utterance alternatives. The English **dative alternation** is a prime example of a locus of choice speakers face, in this case how to order a **theme** and a **recipient**:

- (1) The teacher gave **the very long book** to **the student**.
- (2) The teacher gave **the student** **the very long book**.

We refer to theme-first sentences like (1) as **prepositional object** (PO) constructions, and to recipient-first sentences like (2) as **double object** (DO) constructions. The dative alternation is just one of a much larger class of English syntactic alternations, but has received considerable theoretical interest (for example, Green, 1974; Oehrle, 1976; Gropen, Pinker, Hollander, Goldberg, & Wilson, 1989; Goldberg, 1992; Wasow, 2002; Bresnan, Cueni, Nikitina, & Baayen, 2007; Hovav & Levin, 2008; Slevc, 2011; Futrell, 2023, *inter alia*).

A focus of theories of language production has been to disentangle whether production choices are the result of mere speaker-internal production pressures to produce more mentally **accessible** material early and ease the cost of production (Availability-Based Production, V. S. Ferreira & Dell, 2000; Bock & Warren, 1985), or the result of **meaning-driven** pressures geared towards listeners correctly inferring the speaker's intended meaning (see for example, Lockridge & Brennan, 2002; Brennan & Clark, 1996)

Availability-Based Production has received ample empirical support, including in the context of the dative alternation

(Bock & Irwin, 1980; Bock, 1986; V. S. Ferreira & Griffin, 2003; Vogels, Krahmer, & Maes, 2015; D'Elia, 2016; Koranda, Zettersten, & MacDonald, 2022, *inter alia*). For example, shorter constituents tend to be more accessible than longer ones, so speakers are likely to produce (2) than (1), as **the student** is more accessible than **the very long book**.

One type of meaning pressure is the pressure for **informative** communication (Grice, 1975). An increasing body of work has shown that informativity pressures can explain many pragmatic phenomena at the level of entire sentences or phrases (see Goodman & Frank, 2016; Degen, 2023, for overviews). However, it is well established that production is an incremental process (Bock, 1982; F. Ferreira & Swets, 2002; Brown-Schmidt & Hanna, 2011). The extent to which informativity is taken into account by speakers in incremental production is under-explored. Recent incrementalized versions of pragmatically-oriented theories of communication like the **Rational Speech Act** (RSA) framework predict that speakers produce more informative constituents earlier in an utterance in order to maximize information provided to a listener (Cohn-Gordon, Goodman, & Potts, 2019, 2018).

This prediction stands in direct contrast to the prediction of a different type of accessibility-based theory coming from the literature on information structure: this work has extensively documented a *given before new* preference, explained by a tendency for less informative constituents to be produced earlier because they are more accessible to the speaker (Chafe, 1976; Bock, 1977; Smyth, Prideaux, & Hogan, 1979; Arnold, Losongco, Wasow, & Ginstrom, 2000).

Still other meaning-based accounts posit that informativity plays little to no role in the dative alternation; instead, other constraints on the lexical semantics of the verb (Hovav & Levin, 2008) or on the accessibility of constituents (such as animacy and length) are more influential (Bresnan et al., 2007; Bresnan & Ford, 2010).

Here, we investigate the effects of both informativity and accessibility on the choice of alternant in the dative alternation. We first ask whether incremental constituent informativity affects the choice, using the format of a free spoken production interactive reference game. Players are paired with partners and tasked with describing a target image from a set of four, and we explicitly modulate the relative informativity of the target recipient. To foreshadow, we find no evidence for an effect of informativity. The remainder of the paper thus

studies the possible underpinnings of previously documented accessibility effects in more detail. In particular, we propose that accessibility effects can be characterized through an **automatic production policy** (Futrell, 2023). We show that such a model accurately captures accessibility preferences of constituent length under conservative, empirically supported assumptions on working memory.

Theories of Alternation Preferences

From a semantic perspective, the dative alternation is broadly analyzed in two ways. The **Single Meaning Approach** posits that both the PO and DO constructions are semantically equivalent, and that the alternation is driven entirely by other factors like accessibility and informativity (Baker, 1988; Larson, 1988). In contrast, the **Multiple Meaning Approach** assumes that each construction maintains a distinct meaning and thus a distinct realization of argument structure. This does not necessarily imply that the constructions are truth-conditionally incompatible.

In particular, Hovav and Levin (2008) among others (for example, Gropen et al., 1989; Goldberg, 1992) posit that the PO construction implies caused *motion* of the theme towards the recipient, while the DO construction implies caused *possession* of the theme by the recipient. Hovav and Levin (2008) go on to argue that this is modulated by verb choice: *give* and *sell*, for example, express caused possession, while *throw* and *show* can express caused motion *or* possession.

Evidence for Availability-Based Production

Despite these lexical semantic accounts, it is well attested that there is more variation to be explained: *give*, for example, is extensively used in both the PO and DO construction. Bresnan et al. (2007) show that—in addition to by-verb effects—the dative alternation is also modulated by features of the theme and recipient, including frequency, animacy, constituent length, and givenness.

More generally, these featural preferences are a manifestation of accessibility effects: speakers produce easier-to-retrieve constituents first where syntactically allowed, the main tenet of Availability-Based Production (ABP, V. S. Ferreira & Dell, 2000). Much work has explored various manifestations of these effects, ranging from lexical features like animacy (McDonald, Bock, & Kelly, 1993), to syntactic effects like priming (Bock, 1986), to properties like imageability and visual salience (Bock & Warren, 1985; D’Elia, 2016).

A key prediction of ABP is that referents previously mentioned in discourse are more accessible than (and thus produced before) referents that are *new*, an effect which has seen ample empirical support across domains. This **given-before-new** preference is also a key observation of information structural theories; see (Chafe, 1976; Bock, 1977; Prince, 1981; Arnold et al., 2000; Arnold, Kaiser, Kahn, & Kim, 2013).

Incremental RSA

On the other hand, production choices have also been explained as a result of speakers choosing utterances that are

maximally informative to listeners. This argument goes back to Grice’s (1975) maxims of Quantity and Relation, which posit that listeners expect speakers to provide exactly as much information as is relevant. Models utilizing this kind of argument have especially seen success in modeling pragmatic phenomena such as the production of referring expressions (Frank & Goodman, 2012; Degen, Hawkins, Graf, Kreiss, & Goodman, 2020; Waldon & Degen, 2021).

One way of unifying informativity and accessibility effects is by treating production as a *tradeoff* between accessibility and informativity. For example, Koranda et al. (2022) show that the production of individual lexical items follows a ‘good-enough’ principle balancing informativity with item frequency. For example, we might produce the (more frequent, hence more accessible) word *cat* over *kitten* if *cat* is contextually informative enough to a listener.

Much work has shown that the RSA framework (Frank & Goodman, 2012) is a powerful model for understanding language use in terms of this tradeoff. In RSA, a speaker recursively reasons about a listener in producing an utterance. More formally, a speaker S_1 is a distribution over utterances u conditional on world states s , such that the probability of producing u is proportional to the probability that a ‘literal listener’ $L_0(s | u) \propto [[u]](s)$ would arrive at the intended world state s when given the utterance:

$$S_1(u | s) \propto \exp(\alpha(\log L_0(s | u) - C(u))), \quad (1)$$

where $C(u)$ is a cost function, a term which accounts for accessibility effects.

In other words, a speaker maximizes $L_0(s | u)$ by choosing an utterance that is minimally compatible with the desired state s . For example, suppose a speaker is disambiguating between four possible dative world states, each with the same theme but a unique recipient. Here, the recipient of the target state is the most informative feature to mention to a literal listener, as $L_0(s | \text{recipient}) = 1$, while $L_0(s | \text{theme}) = 1/4$.

RSA has enjoyed empirical success as a model of the production of utterances on a ‘global’ level, e.g. entire constituents or sentences. Here, we ask whether an *incrementalized* version of the RSA, originally proposed by Cohn-Gordon et al. (2019), can explain the production of the dative alternation. In this incremental RSA model, speakers produce utterances word-by-word, maximizing reward at each step. This means that more informative words are predicted to be placed earlier in a sentence – unless very costly/inaccessible, in line with the predictions of ABP and information structural accounts.

Experiment 1: Informativity

In this interactive web-based spoken production experiment, we tested whether informativity affects incremental language production, in particular the choice between DO and PO constructions. We did so by explicitly modulating the relative informativity of the recipient compared to the theme.




Figure 1: Example stimuli for each informativity condition. In the control condition, each image has a unique theme and a unique recipient. In the low informativity condition, each image has one of two themes and a unique recipient. Thus, the target recipient is 2 times more informative than the theme. In the high informativity condition, each image has the same theme and a unique recipient. Thus, the target recipient is 4 times more informative than the theme.

Methods

Participants. We recruited 162 participants on Prolific, yielding 81 dyads paired via Empirica, a virtual lab platform. Of these, 30 dyads (60 participants) were excluded due to issues with recording audio or completing the study.

Materials and procedure. Participants were paired as a **director** and a **guesser**, and placed into an audio call for the duration of the experiment. We used 100ms for audio calls, and Faster Whisper (an implementation of Radford et al.’s (2023) Whisper) for transcription.

Prior to the main experiment phase, participants completed a noun training and a noun recall phase to familiarize them with common-ground labels for the nouns. After the recall phase, they completed a practice phase with two example filler stimuli. In order to control for variation in accessibility, all recipients were animate, and all themes were inanimate. Participants learned single-word labels for each noun, and saw each noun two times before the experiment during the training and recall phase.

On each trial, both participants saw a set of four (shuffled) images displaying actions. The director saw a square placed around the target image and was told to describe the target image to the guesser using a specific target verb shown. The

guesser was asked to select the image corresponding to their partner’s description. After each trial, participants received feedback on whether or not the guesser selected the target.

Each image quartet had the same agent. We manipulated the informativity of the recipient by changing the number of unique recipients and themes (see Fig. 1). This yielded a high and low recipient informativity condition. In the high informativity condition, each image had the same theme and a unique recipient; in the low informativity condition, each image had one of two themes and a unique recipient. We also included a control condition where both the theme and recipient were unique and thus equally informative. Stimuli were modified from Slevc (2011).

We included six critical dative verbs: *give*, *offer*, *sell*, *show*, *throw*, and *toss*. Each dyad encountered each critical verb once and also completed six intransitive and six transitive filler trials, for a total of 18 trials, in random order.

Exclusions We excluded individual trials if the director did not provide a plausible label for each of the referents. For example, we accepted trials where ‘monk’ was mislabeled as ‘priest,’ but not as ‘someone’ (20 trials). In total, we excluded 26.5% of trials. This left 225 total utterances for analysis.

Results and Discussion


Figure 2(a) shows the proportion of DO and PO utterances by informativity condition, alongside incremental RSA model predictions in Fig. 2(b).¹ There was a strong preference towards the PO construction, which made up 84% of utterances. To test whether informativity affected form choice, we conducted a Bayesian mixed-effects logistic regression predicting DO mention from a fixed effect of informativity (dummy-coded, reference level: ‘low informativity’) and random by-item, by-participant, and by-verb intercepts.²

¹We set $\alpha = 1$, and assume uniform cost $C(u) = 0$ of utterances.


²We did not observe evidence of priming, as tested in by-participant logistic regressions predicting the produced order based

	Estimate	1-95% CI	u-95% CI
Intercept	-8.22	-13.10	-4.49
give	4.43	0.69	8.92
offer	5.99	2.28	10.71
sell	2.99	-0.65	7.41
show	5.98	2.36	10.76
toss	2.76	-1.01	7.13

Table 1: By-verb Bayesian mixed effects logistic regression, with *throw* as the reference level.



(a) Empirical preferences. Error bars represent 95% CI.



(b) Incremental RSA predictions. Production probability refers to S_1 .

Figure 2: Empirical and predicted preferences for DO/PO utterances for each level of informativity.

There was no evidence that the control condition resulted in different DO realizations than the low informativity condition ($\beta = 0.15$, $CrI = [-2.00, 2.28]$). There was moderate to strong evidence that the high informativity condition resulted in more DO realizations than the low informativity condition ($\beta = 1.29$, $CrI = [-0.68, 3.41]$, Bayes factor = 10.09), in line with incremental RSA predictions.

However, the magnitude of the effect predicted by RSA is far greater than what we see in the empirical data. Moreover, participants had clear baseline preferences that were not captured by the RSA model; in particular, across the board the PO order was preferred to the DO order. One way of mitigating this is to equip the RSA model with a heavily PO-biased prior. It's also possible, however, that incremental informativity considerations play at most a minor role in the choice of alternant. If so, an explanation for this might be that speakers do not incrementally optimize (much) for listeners because such optimization is computationally expensive. Instead, speakers might optimize for listeners at the utterance level, while incremental production is for the most part influenced by factors like accessibility and verb semantics.

An alternative explanation is that speakers *do* incrementally optimize, but accessibility effects are so pronounced in the relatively constrained case of the dative alternation that we do not see these effects borne out. It would also be valuable to test informativity effects in a less constrained setting. The utterances evaluated here were short, consisting of only four constituents. Longer utterances might yield greater informativity effects on production preferences.

The collected dataset also allows us to test predictions made by Hovav and Levin's (2008) analysis of the Multiple Meaning Approach, under which different verbs should prefer different amounts of DO realization. Proportions of DO realizations by verb are shown in Fig. 3. To test for by-verb variation, we ran a Bayesian mixed effects logistic regression predicting DO mention from a fixed effect of verb (dummy-

on trial number. None of these met the threshold for significance ($p > 0.1$ for all).

coded, reference level: 'throw'), with random by-item and by-participant intercepts. Results are in Table 1. Verb signifi-




Figure 3: Proportion of DO/PO utterances for each verb, aggregate over all informativity conditions. Error bars represent bootstrapped 95% confidence intervals.




Figure 4: Proportion of DO/PO utterances in the Switchboard dataset analyzed by Bresnan et al. (2007) for verbs occurring in our Exp. 1.

cantly impacted the realization of the dative alternation, with *give*, *offer*, and *show* being more frequently realized as DO than *throw*, *toss*, and *sell*.

This is consistent with Hovav and Levin’s (2008) analysis. In particular, *throw* and *toss* (expressing caused motion) were almost exclusively realized as PO, while *give* and *offer* (expressing possession) took the DO realization more frequently.

However, we still observed an overwhelming preference towards the PO construction across verbs; note that this result notably differs from that of Bresnan et al. (2007), who found a much greater preference for DO realizations (see Fig. 4). We posit that this is an effect of the stimuli we used: in each image, the agent is currently in possession of the theme, and there is an expectation that the theme will be moved *towards* the recipient. In effect, we coerce a caused motion reading of each verb, consistent with Hovav and Levin’s (2008) analysis. In contrast, Bresnan et al.’s (2007) analysis used the Switchboard Corpus, which consists of naturally occurring speech.

Experiment 2: Accessibility

In Exp. 1, we found that informativity alone is likely not enough to explain the dative alternation at an incremental level. However, models like incremental RSA express trade-offs between both informativity- and accessibility-driven effects. Here, we shift focus to the predictions of Availability-Based Production, aiming to more rigorously characterize the notion of incremental accessibility.

Most computational models of production, including RSA, eschew a concrete source of accessibility effects, instead fitting an arbitrary cost function to data based off of certain featural preferences like frequency and animacy. This is theoretically unsatisfying, as it provides no insight as to *where* accessibility preferences originate.

Slevc (2011) shows that accessibility effects may arise in part due to working memory effects. In particular, he demonstrates that accessibility effects decrease in the presence of working memory load, as memory load interferes with the accessibility of a referent; this has parallels to work on similarity-based interference in the literature of language processing (Lewis, 1996; Oberauer, 2002; Lewis, Vasishth, & Van Dyke, 2006; Timkey & Linzen, 2023).

Further, Arnold et al. (2013) argue that information structural phenomena, including the given-before-new preference, can also arise from memory effects: ‘given’ information is that which is already present in working memory, while ‘new’ information has yet to be retrieved from long-term memory.

Futrell (2023) proposes that incremental language production is in part modulated by an **automatic policy** p_0 , which is a distribution over utterances u marginalized over all possible world states. That is, we can think of p_0 as an autoregressive language model in the natural language processing sense. In essence, this replaces the cost function $C(u)$ of RSA.

Intuitively, more accessible utterances tend to be more predictable from context. For example, the given-before-new pattern straightforwardly falls out of this policy, since given

constituents are generally more contextually predictable. However, Futrell (2023) does not rigorously show how the automatic policy might give rise to specific accessibility effects.³ Here, we expand upon this line of work, concentrating on the **short-before-long** constituent length preference:

- (3) ?The teacher gave the very long hardcover book by a famous author to the student.
- (4) The teacher gave the student the very long hardcover book by a famous author.

Here, the DO construction is more acceptable than the PO variant (Stallings & MacDonald, 2011). We demonstrate that this effect can be explained via Futrell’s (2023) automatic policy when equipped with a limited working memory.

More formally, we define a language model as a distribution p over utterances u given context c . Rather than predicting u based off of the true context c , we give the language model a **lossy representation** of the context $M(c) = m$, which gives the model only partial information about the true context. This lossy context assumption has seen broad success in the language comprehension literature as a way to unify expectation-based and memory-based models and explain a variety of online processing effects (Futrell, Gibson, & Levy, 2020; Hahn, Futrell, Levy, & Gibson, 2022).

A lossy-memory language model must make predictions about u based on partial information such that the longer the intervening context between u and the start of the sentence, the more information is lost. Thus, we would expect such a model to prefer placing shorter constituents earlier. This argument is similar to the interference-based argument of Slevc (2011): longer constituents interfere with retrieval of context.


Methods

We trained language models with a naïve loss model: **fixed-length context windows**. The model makes predictions based only on the previous n tokens. This is a very straightforward implementation of lossy memory which can be improved upon (for example, see Futrell et al., 2020; Hahn et al., 2022); this induces a memory bottleneck at the level of short-term information storage. We empirically evaluated the performance of both n -gram models and Long-Short Term Memory RNNs (LSTMs) with fixed context windows.


We modulated the size of the context window from 4 to 7 for n -grams and from 3 to 6 for LSTMs. However, if the modifier is too short for a given context window size, the model is effectively lossless. Therefore, we additionally modulated the length of each constituent in the training and test sets between 2 and 5 tokens. In total, we trained 16 models for each architecture.

We aim to determine whether these models can learn a short-before-long preference solely based on limited memory capacity. Thus, we trained them on a ‘balanced’ toy dataset

³The automatic policy predicts accessibility effects if these effects are present in source distribution. However, to work out the *source* of accessibility, we consider a data-agnostic setting.



(a) n -gram



(b) LSTM

Figure 5: Average preference for PO order by theme and recipient length for models with context size n and modifier length k . Error bars represent 95% CI. Plots marked with significance of t -test between conditions. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

of dative structures. We generated tuples of two nouns and two modifiers which are uniformly drawn from a simulated vocabulary, and then generated all sixteen possible datives, expressing short-short, short-long, long-short, and long-long preferences for each of the four noun-modifier pairings. This ensured that there were no effects by lexical item, and that short- and long-first orders were equally present in the training data. The training set consisted of 64000 such sentences; the test set consisted of another 8000.

We then measured the preference for producing the PO order using the **surprisal** of each construction. Formally, the surprisal of a sequence of tokens $w_1, \dots, w_n$ is

$$s(w_1, \dots, w_n) = -\log p(w_1, \dots, w_n). \quad (2)$$

We then measured $s(\text{DO construction}) - s(\text{PO construction})$, which will be positive when the PO order is *lower* surprisal (i.e. more predictable). That is, we expect a positive value when the recipient is long (and the theme is short), and a negative value when the theme is long (and the recipient is short).

Results

Fig. 5 shows the results of the n -gram and LSTM models. In both models, we observed effects consistent with the short-before-long preference, as long as the context window is not so long that the context becomes lossless.

In particular, the PO order was preferred in the short theme, long recipient case, while the DO order was preferred in the long theme, short recipient case across conditions. We found a significant effect in preference across all relevant conditions by two sample t -test.

In short, we showed that constituent length preferences fall out of an autoregressive language model with basic restrictions on working memory capacity. This is equivalent to Futrell’s (2023) automatic policy with lossy context, providing compelling evidence towards characterizing the source of accessibility effects in language production.

Conclusion

Here, we studied informativity and accessibility effects in incremental production of the dative alternation. We demonstrated that while referent informativity does have a slight effect on choice of alternant, it is far smaller than the magnitude predicted by incremental RSA. We then proposed a model of accessibility effects based on principles of lossy memory, demonstrating that such a model induces a short-before-long constituent length preference. Our findings help characterize both meaning-based and accessibility driven effects.

Much past work has found that informativity is an important pressure for production of entire utterances on the global level (Degen, 2023). Our findings provide an interesting wrinkle in this narrative: such ‘macro-level’ principles of production do not directly translate to incremental word-by-word production. It is an open question as to why this is—perhaps incremental informativity is computationally expensive, or accessibility effects are simply more powerful at the incremental level. An alternative possibility is that in our setup of the task in Exp. 1, participants were not sensitive to differences between images, fixating on the target image rather than processing differences between the individual scenes. Future work should address this issue.

In contrast, we are able to more formally characterize accessibility in terms of Futrell’s (2023) automatic policy. This provides a more rigorous basis for the origin of features that drive ABP and good-enough production, stepping away from the black box assumption of much previous work, and serves an explanatory tool for principles of information structure. It remains to be seen if this characterization can explain a wider variety of empirically observed effects of accessibility. This adds to a growing body of work suggesting that working memory processes can explain effects of online language production and processing (Futrell et al., 2020; Merks & Frank, 2020; Schrimpf et al., 2021; Hahn et al., 2022; Goldberg & Ferreira, 2022; Timkey & Linzen, 2023).

Acknowledgments

This work benefited from discussion at BlackboxNLP 2023 and CAMP[6] 2024. We also thank Emily Goodwin and Richard Futrell for their helpful feedback and comments. All code and data is available at <https://github.com/neilrathi/ditransitive>.

References

- Arnold, J. E., Kaiser, E., Kahn, J. M., & Kim, L. K. (2013). Information structure: linguistic, cognitive, and processing approaches. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(4), 403–413.
- Arnold, J. E., Losongco, A., Wasow, T., & Ginstrom, R. (2000). Heaviness vs. newness: The effects of structural complexity and discourse status on constituent ordering. *Language*, 76(1), 28–55.
- Baker, M. C. (1988). *A theory of grammatical function changing*. Doctoral, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Bock, J. K. (1977). The effect of a pragmatic presupposition on syntactic structure in question answering. *Journal of Verbal Learning and Verbal Behavior*, 16(6), 723–734.
- Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological review*, 89(1), 1.
- Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, 18(3), 355–387.
- Bock, J. K., & Irwin, D. E. (1980). Syntactic effects of information availability in sentence production. *Journal of Verbal Learning and Verbal Behavior*, 19(4), 467–484.
- Bock, J. K., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, 21(1), 47–67.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of experimental psychology: Learning, memory, and cognition*, 22(6), 1482.
- Bresnan, J., Cueni, A., Nikitina, T., & Baayen, R. H. (2007). Predicting the dative alternation. In *Cognitive foundations of interpretation* (pp. 69–94). KNAW.
- Bresnan, J., & Ford, M. (2010). Predicting syntax: Processing dative constructions in american and australian varieties of english. *Language*, 86(1), 168–213.
- Brown-Schmidt, S., & Hanna, J. E. (2011). Talking in another person's shoes: Incremental perspective-taking in language processing. *Dialogue & Discourse*, 2(1), 11–33.
- Chafe, W. (1976). Givenness, contrastiveness, definiteness, subjects, topics, and point of view. *Subject and Topic*.
- Cohn-Gordon, R., Goodman, N., & Potts, C. (2018, June). Pragmatically informative image captioning with character-level inference. In M. Walker, H. Ji, & A. Stent (Eds.), *Proceedings of the 2018 conference of the north American chapter of the association for computational linguistics: Human language technologies, volume 2 (short papers)* (pp. 439–443). New Orleans, Louisiana: Association for Computational Linguistics. doi: 10.18653/v1/N18-2070
- Cohn-Gordon, R., Goodman, N., & Potts, C. (2019). An incremental iterated response model of pragmatics. In G. Jarosz, M. Nelson, B. O'Connor, & J. Pater (Eds.), *Proceedings of the society for computation in linguistics (SCiL) 2019* (pp. 81–90). doi: 10.7275/cprc-8x17
- Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, 9, 519–540.
- Degen, J., Hawkins, R. D., Graf, C., Kreiss, E., & Goodman, N. D. (2020). When redundancy is useful: A bayesian approach to “overinformative” referring expressions. *Psychological Review*, 127(4), 591.
- D'Elia, S. C. (2016). *The spray/load and dative alternations: Aligning vp structure and contextual effects*. Doctoral, University of Kent, Kent, United Kingdom.
- Ferreira, F., & Swets, B. (2002). How incremental is language production? evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language*, 46(1), 57–84.
- Ferreira, V. S., & Dell, G. S. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive psychology*, 40(4), 296–340.
- Ferreira, V. S., & Griffin, Z. M. (2003). Phonological influences on lexical (mis) selection. *Psychological Science*, 14(1), 86–90.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Futrell, R. (2023). Information-theoretic principles in incremental language production. *Proceedings of the National Academy of Sciences*, 120(39), e2220593120.
- Futrell, R., Gibson, E., & Levy, R. P. (2020). Lossy-context surprisal: An information-theoretic model of memory effects in sentence processing. *Cognitive Science*, 44(3), e12814.
- Goldberg, A. E. (1992). The inherent semantics of argument structure: The case of the English ditransitive construction. *Cognitive Linguistics*, 3(1), 37–74. doi: 10.1515/cogl.1992.3.1.37
- Goldberg, A. E., & Ferreira, F. (2022). Good-enough language production. *Trends in Cognitive Sciences*, 26(4), 300–311.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
- Green, G. M. (1974). *Semantics and syntactic regularity*. Bloomington: Indiana University Press.
- Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41–58). Brill.
- Gropen, J., Pinker, S., Hollander, M., Goldberg, R., & Wilson, R. (1989). The learnability and acquisition of the dative alternation in english. *Language*, 203–257.
- Hahn, M., Futrell, R., Levy, R., & Gibson, E. (2022). A resource-rational model of human processing of recursive

- linguistic structure. *Proceedings of the National Academy of Sciences*, 119(43), e2122602119.
- Hovav, M. R., & Levin, B. (2008). The English dative alternation: The case for verb sensitivity1. *Journal of Linguistics*, 44(1), 129–167.
- Koranda, M. J., Zettersten, M., & MacDonald, M. C. (2022). Good-enough production: Selecting easier words instead of more accurate ones. *Psychological Science*, 33(9), 1440–1451.
- Larson, R. K. (1988). On the double object construction. *Linguistic inquiry*, 19(3), 335–391.
- Lewis, R. L. (1996). Interference in short-term memory: The magical number two (or three) in sentence processing. *Journal of psycholinguistic research*, 25(1), 93–115.
- Lewis, R. L., Vasishth, S., & Van Dyke, J. A. (2006). Computational principles of working memory in sentence comprehension. *Trends in cognitive sciences*, 10(10), 447–454.
- Lockridge, C. B., & Brennan, S. E. (2002). Addressees' needs influence speakers' early syntactic choices. *Psychonomic bulletin & review*, 9(3), 550–557.
- McDonald, J. L., Bock, K., & Kelly, M. H. (1993). Word and world order: Semantic, phonological, and metrical determinants of serial position. *Cognitive psychology*, 25(2), 188–230.
- Merkx, D., & Frank, S. L. (2020). Human sentence processing: Recurrence or attention? *arXiv preprint arXiv:2005.09471*.
- Oberauer, K. (2002). Access to information in working memory: exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 411.
- Oehrle, R. T. (1976). *The grammatical status of the English dative alternation*. Doctoral, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Prince, E. F. (1981). Toward a taxonomy of given-new information. *Radical Pragmatics*.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2023). Robust speech recognition via large-scale weak supervision. In *International conference on machine learning* (pp. 28492–28518).
- Schrimpf, M., Blank, I. A., Tuckute, G., Kauf, C., Hosseini, E. A., Kanwisher, N., ... Fedorenko, E. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45), e2105646118.
- Slevc, L. R. (2011). Saying what's on your mind: Working memory effects on sentence production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(6), 1503.
- Smyth, R. H., Prideaux, G. D., & Hogan, J. T. (1979). The effect of context on dative position. *Lingua*, 47(1), 27–42.
- Stallings, L. M., & MacDonald, M. C. (2011). It's not just the "heavy np": relative phrase length modulates the production of heavy-np shift. *Journal of Psycholinguistic Research*, 40, 177–187.
- Timkey, W., & Linzen, T. (2023). A language model with limited memory capacity captures interference in human sentence processing. *arXiv preprint arXiv:2310.16142*.
- Vogels, J., Krahmer, E., & Maes, A. (2015). How cognitive load influences speakers' choice of referring expressions. *Cognitive science*, 39(6), 1396–1418.
- Waldon, B., & Degen, J. (2021). Modeling cross-linguistic production of referring expressions. In *Proceedings of the society for computation in linguistics 2021* (pp. 206–215).
- Wasow, T. (2002). *Postverbal behavior*. CSLI Publications.