

Probabilistic Pragmatics and Probabilistic Experience

AMIR ZELDES
GEORGETOWN UNIVERSITY
amir.zeldes@georgetown.edu

In their article about Probabilistic Pragmatics, Franke and Jäger (F&J) raise important issues for some of the tasks facing pragmatics, and how progress can be made: rather than only noting the possibilities for the expression of meaning under pragmatic modulation, or coding them in a formal algebraic manner, we gain a more precise description and a computationally implementable predictive model by incorporating probabilities and Bayesian reasoning into our models of speakers' behavior (Frank & Goodman 2014, Franke 2014; see also Lassiter, to appear, on Bayesian semantics). The explanatory insight that we can get into why speakers make gradient inferences, or use strategies such as indirect speech acts, falls out of the Bayesian approach naturally. In parallel to this, there is another level of probabilistic description that came to my mind as I read their paper, which I feel should also be emphasized if we are aiming at predictive power with regard to speakers' behavior (i.e. answering the question, '*what will speaker S say next, given that..?*'). That level exists in the dimension of experience, or more precisely in terms of constructionist and cognitive approaches, entrenchment and prototypicality in our knowledge of the usage of our language (Langacker 2000, Bybee 2006, Goldberg 2006, among others). This level of description is informed not only by experiments of the type discussed by F&J, but also to a large extent by consulting naturalistic corpus data, which is mentioned only in passing in the article.

Over a decade ago, Manning (2003) presented a research agenda for syntactic theory which he termed 'Probabilistic Syntax'. A main point of his approach, which is echoed throughout multifactorial work in syntax (Gries 2002, Bresnan et al. 2007, Lohmann 2011), was the goal of modeling speaker output given a meaning to be expressed and an appropriate context, formulated as a conditional probability in (1), and conversely, the ability to infer meaning given a linguistic form and its context, as in (2):

- (1) $p(\text{form}|\text{meaning, context})$
- (2) $p(\text{meaning}|\text{form, context})$

The task of the speaker is to choose an appropriate form given the meaning they wish to express and the context, and the task of the hearer is to infer the correct meaning given the form and the context.¹ The probabilities are inferred not just by pragmatically motivated cost based estimates, optimizing for the rational choices of each participant: they are modelled, among other things, on previous experience.

As an example of how these two sources of probabilistic information complement each other, let us consider the case of indirect speech acts. F&J point out that it should be most conducive to getting another glass of wine to say directly “Pour me another glass of wine!” (p. 25), though clearly, in many situations, it is a dispreferred form of expression. Predicting what speakers will actually say can only be discussed in the context of a competition between multiple options. This brings with it the challenge at the outset of any Variationist study: determining the variable or ‘variant field’ to be explored and the variants which instantiate it (see Labov 2004, Rissanen 2008). Once the variants to be considered are agreed upon, we can associate different probabilities with each outcome and speak of factors co-determining the decision. In phonology, this space is more limited categorically (though limitless on an analogue scale), so that Labov, discussing conditions for the English copula contraction and deletion, argued already in the late sixties (1969:737-738) that a discrete set of outcomes could be modeled such that we “associate with each VARIABLE RULE a specific quantity ϕ which denotes the proportion of cases in which the rule applies as a part of the rule structure itself”. What rule motivates an indirect speech act, what are the relevant variants, what are their quantities and which predictors can help us to model the response?

For a variable like ‘requesting something’ there can be countless possible realizations. In the interest of limiting the scope of this example, I will consider scenarios similar to Franke and Jäger’s request of a glass of wine: a first person, present-anchored factual request for something (as opposed to hypotheticals, narratives about others’ requests, etc.). To operationalize this, I will be looking at first person subjects in spoken data using four variants: *Give me ...!*, *I want*, *I’d like*

¹ We could also add that the task of second language learners is to learn these mappings on top of the native ones they are equipped with (see Kasper & Rose 2003), and that the task of Natural Language Understanding and Generation systems is to model precisely these probabilities as well.

and *I wouldn't mind* (though countless others are possible), followed by an indefinite object.² Of these options, only the first is truly 'direct', but I gauge 'want' as most direct after that (present indicative) and the rest in decreasing order (modal 'would like', and most plausibly deniable 'wouldn't mind') – these seem to progressively 'downplay interest', following F&J. From a pragmatic point of view, I would expect less direct approaches to correlate with more face-threatening situations, though the extreme cases of embarrassment due to attempted bribery discussed by F&J cannot be found easily in spontaneous corpus data.

We can examine the relative frequencies of these variants in spoken (British) English using data from the spoken section of the British National Corpus (BNC, <http://www.natcorp.ox.ac.uk/>, approx. 10 million tokens of spoken data). Although any corpus we select will only be useful in predicting the behavior of 'more data of the same kind', the BNC is uniquely suitable here in providing a very large component of transcribed spontaneous conversations, which lend themselves to authentic interpersonal requests. A surface search for the forms with subject *I* and a marked indefinite object or a clear, unambiguous sentence initial *give*, was filtered manually for actual interpersonal requests. Excluded sentences that are not construed as requests include wishful expressions of desire (not targeting the hearer) as in (3), reported speech as in (4), or requests for the non-immediate future (5):³

(3) *I'm moving into a new flat soon and I want a living room, a lounge, a lounge, bedroom*

(KB7: not a request that the listener is expected to fulfill)

(4) *she said to me Sue I want a lot, all over.* (KB6: reported speech, not necessarily verbatim)

(5) - *See you Mandy.* - *Bye.* - *Give me a ring.* - *Yeah*

(KGU: non-immediate, tentative request for a future phone call)

Referring to the issue of indirect speech acts, Manning (2003:291) writes "[w]hen someone says "it's cold in here", in some circumstances I'm understanding them correctly if I interpret that utterance as a request to close the window". What are the circumstances that condition the forms above? Based on the BNC data, the most common way to request something

² I avoid lexical imperatives beyond general 'give', since items like 'pour' are sparse and hard to enumerate exhaustively; a similar study of such forms is of course conceivable. BNC document identifiers are given after each example.

³ The classification of requests and variants undertaken here is not meant to be complete – a study going beyond the present illustrative purpose would have to devise and test explicit guidelines, measure inter-annotator agreement etc.

such as a glass of wine in spoken British English from among these options is the one with ‘want’, though recall that only sentence initial imperative ‘give’ is considered for that verb (and not: ‘could you give..?’). The other forms are substantially less frequent. The form with ‘want’ is not only the most common in interpersonal requests, it is also most frequent in total if meaning is ignored (see Figure 1).




Figure 1. Frequencies in the spoken part of the BNC for *Give me..!*, *I want*, *would like* and *wouldn't mind*, broken down by requests, address to family members and soliciting food.

The other options, though much less frequent, have a higher cue validity (cf. Ellis 2006, Gries 2013): if we encounter one of the other options, we can be more certain that the speaker intended to express an actual request (see the relative height of the left two columns of each group). The data for the ‘wouldn’t mind’ option referred to by F&J is too rare to draw any quantitative conclusions, but the difference for ‘want’ versus ‘give’ and ‘would like’ is fairly pronounced at 64% actual requests for ‘like’ and 58% for give, but only 41% for ‘want’.

While simple frequency is important (prefer the most frequent expression for a function, all other things being equal) a lot of the cues that govern which lexical choices get realized are contextual: these include pragmatics, but also recency or dispersion effects, entropy in terms of information content, i.e. the predictability of a word given the immediately preceding context, and neighboring family size, i.e. the variety of lexical types that are known to occur adjacently to the position of a lexical choice (Baayen 2010).

From a manual inspection of the BNC result, it seems that part of the cues that correlate with the more direct ‘want’ variant are related to the setting and speaker relations: ‘want’ requests are very often uttered between identifiable family members (based on inspection of immediate context, e.g. including vocative ‘mum’), indeed most often from child to mother (6-7), and also exhibit a particularly high likelihood of the object being an item of food (6) or drink (8) (these two tendencies are of course not orthogonal).

(6) *Mummy I want a big bit* (KD1: in reference to meat)

(7) *Mum mum mummy I want a straw* (KBW)

(8) *I want a cup of tea* (4 times: KB6, KBD and twice in KPE)

The more polite ‘would like’ has only 18% requests for food and 14% requests to family members, while ‘want’, as expected for a less polite or indirect construction, refers to food in 55% of requests and is addressed to family in 31% of requests. ‘Give’ is a very broad verb, and has only a handful of occurrences in these two conditions, but this may be a result of its literal meaning: for food ‘give’ could mean handing an item over, not necessarily preparing some food for the speaker, as in (9) versus (10a). If the sandwich has not been made yet, a sentence like (10b) to replace the middle utterance might be dispreferred (marked #?) if the implication is that the hearer should make the food, since *give!* would issue an order for substantially more work than simply handing over a mint. (10b) might be more appropriate when speaking to a waiter.

(9) *Give me a mint cos I haven’t* (KDA)

(10) a. - *Go cut you some bread?*

- *I’m alright, I’d I want a butty thank you.* [a type of sandwich -AZ]

- *Michael how, how many rounds do you want?* (KD8)

b. #? *I’m alright, give me a butty*

Beyond the best fit for lexical semantic meaning (handing over/preparing/...), it is likely that pragmatically motivated and conventionalized cues, such as face-threatening situations, formality and typical choice of objects could all be harnessed for a multifactorial regression model predicting constructional choice in the vein of F&J’s discussion, though clearly more

corpus data would be useful in characterizing the differences between the constructions, and some factors may turn out to be co-linear or redundant with others.

What can we learn about pragmatics from the distribution of the forms? Some factors are likely to be central to human interaction and might not be affected by training or linguistic experience – it is unlikely that we can raise an ‘anti-Grice’ child. For example, speaker familiarity, reputation-building scenarios and risk estimation play an important part in the tendency for indirect speech acts. However it is still difficult to predict which variant will be used in which instance: ‘wouldn’t mind’ appears to be very rare in the spoken BNC data, and ‘would like’, which is also not a literal request, is still less frequent than ‘want’ as a form of request. Seen from a historical perspective, all three forms apart from ‘give’ were etymologically even more indirect compositional expressions that are now entrenched and conventionalized: *want* originally meant ‘lack’ (cf. *for want of a nail ... the kingdom was lost*) and *like* meant ‘to fit’ (originally with oblique experiencer, *me likes it*, i.e. ‘it fits, pleases me’, later active *I like*).⁴ In either case, the speaker is mentioning that they lack something without requesting it, or that it would be fitting for them to have it, again with no direct request or even reference to their desiring anything. That these became increasingly interpreted as desiderative, more direct forms of request, especially in the case of *want*, is a testament to the strong functional pressure to realize requests as indirect speech acts; but also to the power of grammaticalization, resulting from the abductive reasoning which associates the common function of ‘(polite) request’ with a no longer analyzed, by now arbitrary Saussurean *signifiant* (see also Croft 2010 on real world motivation in grammaticalizing frequent expressive strategies).

Indeed, in the puzzling example of Dexter Morgan’s unsuccessful and subsequent successful marriage proposals discussed by F&J we see an echo of the same factors of ‘guilt by association’: Proposal #1 is, in terms of prototypes and entrenched lexical choice an extremely bad proposal, and this is the humor of it. The words *insurance*, *deduction* and *dependent* do not occur once within 5 words of ‘marry’ in the entire BNC, suggesting that they might have the wrong ‘semantic prosody’ (cf. Sinclair 2004:173), i.e. a ‘spreading of connotational colouring beyond single word boundaries’ (Partington 1998:68; see Stewart 2010 for a critical discussion). The words *family*, *feel* and *kids/children* all do (416 times in total), with *family* (the main thrust of Dexter’s proposal) being a very strong collocate (rank 22 among nouns rated by log-

⁴ See "like, v.1." and "want, v.". OED Online. Oxford University Press, June 2015. Web. 31 August 2015.

likelihood, for a 5 token window; *children* is place 7; *kids* is at 79, but is a lot less frequent in British English than Dexter's American English).⁵ Needless to say, there are real world reasons, including pragmatic reasons why the concept of 'family' makes for a better context for a marriage proposal than 'insurance', but as recipients of the constant stream of language input around us, we cannot help but learn the content words, functions, and syntax that make a good proposal sound right (e.g. often conventionalized 'will you marry me' with archaic *will* meaning 'to want', not compositional/spontaneous 'do you want to marry me'), and the same may be true of how we learn to ask people to close the window in our language. I think that there is little doubt that the differences between the usage of 'attractive' and 'stunning' as mentioned by F&J (p.14), whether scalar or not, cannot be addressed without recourse to linguistic experience and distributional data. Studying the interaction of usage knowledge and pragmatics constraints will therefore require both experimental and corpus based study designs.

It is very likely that the ideas presented above can be integrated into the probabilistic framework put forth by F&J as further priors on variant probabilities interacting in the multifactorial model with the pragmatic constraints being modeled. What is less certain is that we can tease apart which motivation increases the weight of a cue in a predictive regression model: is a cue useful due to the rational risk-reducing choice behind, say, an indirect speech act, or is it due to the effect of rote memorization behind the conventional choice of a form corresponding to such indirectness? The former motivation is the reason why such forms arise in the first place; the latter is the reason why we forget, or cease to understand, that what are now conventionalized, rather direct forms of expression were once more circumspect, indirect ones.

References

- Baayen, R. Harald (2010), Demythologizing the Word Frequency Effect: A Discriminative Learning Perspective. *The Mental Lexicon* 5(3), 436–461.
- Bresnan, Joan, Anna Cueni, Tatiana Nikitina & R. Harald Baayen (2007), Predicting the Dative Alternation. In: Gerlof Bouma, Irene Kraemer & Joost Zwarts (eds.), *Cognitive Foundations of Interpretation*. Amsterdam: Royal Netherlands Academy of Science, 69–94.
- Bybee, Joan L. (2006), From Usage to Grammar: The Mind's Response to Repetition. *Language* 82(4), 711–733.
- Croft, William (2010), The Origins of Grammaticalization in the Verbalization of Experience. *Linguistics* 48(1), 1–48.

⁵ A count in the roughly 450 million-word Corpus of Contemporary American English (Davies 2008-) compared to the BNC gives a normalized ratio of 1.4:1 for *children*, but almost 7:1 for *kids*.

- Davies, Mark (2008-), *The Corpus of Contemporary American English: 450 Million Words, 1990-Present*. Available at: <http://corpus.byu.edu/coca/>.
- Ellis, Nick C. (2006), Language Acquisition as Rational Contingency Learning. *Applied Linguistics* 27(1), 1–24.
- Frank, Michael C. & Noah D. Goodman (2014), Inferring Word Meanings by Assuming that Speakers are Informative. *Cognitive Psychology* 75, 80–96.
- Franke, Michael (2014), Typical Use of Quantifiers: A Probabilistic Speaker Model. In: Paul Bello, Marcello Guarini, Marjorie McShane & Brian Scassellati (eds.), *Proceedings of CogSci 2014*. Austin, TX, 487–492.
- Goldberg, Adele E. (2006), *Constructions at Work: The Nature of Generalization in Language*. Oxford: Oxford University Press.
- Gries, Stefan Th. (2002), Evidence in Linguistics: Three Approaches to Genitives in English. In: Ruth M. Brend, William J. Sullivan & Arle R. Lommel (eds.), *LACUS Forum XXVIII: What Constitutes Evidence in Linguistics?* Fullerton, CA: LACUS, 17–31.
- Gries, Stefan Th. (2013), *Statistics for Linguistics with R: A Practical Introduction*. 2nd edition. Berlin: Mouton de Gruyter.
- Kasper, Gabriele & Kenneth R. Rose (2003), *Pragmatic Development in a Second Language*. Oxford: Wiley-Blackwell.
- Labov, William (1969), Contraction, Deletion, and Inherent Variability of the English Copula. *Language* 45(4), 715–762.
- Labov, William (2004), Quantitative Analysis of Linguistic Variation. In: Ulrich Ammon, Norbert Dittmar, Klaus J. Mattheier & Peter Trudgill (eds.), *Sociolinguistics. An International Handbook of the Science of Language and Society*. Vol. 1. Berlin: Walter de Gruyter, 6–21.
- Langacker, Ronald W. (2000), A Dynamic Usage-based Model. In: Michael Barlow & Suzanne Kemmer (eds.), *Usage Based Models of Language*. Stanford: CSLI Publications, 1–63.
- Lassiter, Daniel (to appear), Linguistic and Philosophical Considerations on Bayesian Semantics. In: Matthew Chrisman & Nate Charlow (eds.), *Deontic Modals*. Oxford: Oxford University Press.
- Lohmann, Arne (2011), *Help vs help to*: A Multifactorial, Mixed-Effects Account of Infinitive Marker Omission. *English Language and Linguistics* 15(3), 499–521.
- Manning, Christopher D. (2003), Probabilistic Syntax. In: Rens Bod, Jennifer Hay & Stefanie Jannedy (eds.), *Probabilistic Linguistics*. Cambridge, MA: MIT Press, 289–341.
- Partington, Alan (1998), *Patterns and Meanings: Using Corpora for English Language Research and Teaching*. (Studies in Corpus Linguistics.) Vol. 2. Amsterdam & Philadelphia: John Benjamins.
- Rissanen, Matti (2008), Corpus Linguistics and Historical Linguistics. In: Anke Lüdeling & Merja Kytö (eds.), *Corpus Linguistics. An International Handbook*. Vol. 1. Berlin: Mouton de Gruyter, 53–68.
- Sinclair, John (2004), *Trust the Text*. London and New York: Routledge.
- Stewart, Dominic (2010), *Semantic Prosody. A Critical Evaluation*. New York & London: Routledge.