

This is a comment on a target article by Franke & Jäger, focusing primarily on their discussion of reference games. That article can be found at <http://ling.upenn.edu/~jonsteve/target.pdf>

When do we think strategically?

Jon Scott Stevens*

Michael Franke and Gerhard Jäger (this volume), henceforth F&J, offer a programmatic illustration of the benefits of a Bayesian rational approach to pragmatic analysis. F&J outline three test cases: (i) computational modeling of experimental results involving so-called reference games (Frank and Goodman 2012), (ii) prediction of gradient acceptability judgments of quantifiers, and (iii) a sketch of a theoretical model of indirect speech act production/interpretation in a negotiation scenario. In this comment, I will focus mostly on the first of these test cases, which I believe raises some interesting deeper questions about how to model gradient linguistic behavior, and about the explanatory power of different approaches to cognitive modeling in general. The overall arc of the argument can be summarized as follows.

1. Subjects' responses in F&J's reference game experiment could be interpreted in two ways: First, they could be interpreted as they are by F&J—as the product of some inherently probabilistic approximation of optimal reasoning about speaker intention. Alternatively, they could be interpreted as the interaction of two distinct “winner-takes-all” forms of reasoning. For example, hearer responses in F&J's reference game could be analyzed as the result of an interaction between one mechanism which dumbly zeroes in on a visually salient object and checks whether that object is consistent with the speaker's message, and another (which only kicks in when the first fails) for figuring out which object is strategically optimal under a simple iterated best response (IBR) model (Franke 2009).
2. A simple parameter-free computational model along these lines not only predicts results in the right neighborhood, but also predicts the priors elicited in F&J's prior elicitation experiment with minimal assumptions. Moreover, adding a single free parameter, a probability of defecting from an IBR-based strategy back to a salience-based strategy, results in a near-perfect fit to the observed data.
3. Though both models are only meant to be illustrative, the differences between them highlight the need to consider some important questions: Should we be so quick to abandon “winner-takes-all” models of strategic reasoning, e.g., expected utility-maximizing approaches like IBR, as possible sources of explanation for experimental observations? Might we gain something by separating our models into component parts, representing strategic vs. non-strategic thinking? If so, when do we think strategically? And to what extent can proving the categorical optimality of an action under a normative model serve to explain a statistical tendency observed in the laboratory? These questions are important to this enterprise, and they should not be waved away by appealing to Marr levels (Marr 1982).

These points are taken up in turn in the following three sections.

*Corresponding Author: Jon Stevens, Zentrum für Allgemeine Sprachwissenschaft, Berlin, E-Mail: stevens@zas.gwz-berlin.de

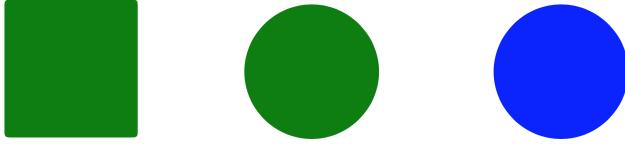


Figure 1: A configuration of shapes for a one-shot reference game experiment

| | GS | GC | BC | m_G | m_S | m_B | m_C |
|------|------|------|------|-------|-------|-------|-------|
| GS | 1,1 | 0,0 | 0,0 | ✓ | ✓ | ✗ | ✗ |
| GC | 0,0 | 1,1 | 0,0 | ✓ | ✗ | ✗ | ✓ |
| BC | 0,0 | 0,0 | 1,1 | ✗ | ✗ | ✓ | ✓ |

Figure 2: Set-up for an IBR derivation of optimal behavior in a reference game.

1 Optimality and gradience

We start with a “pure” game-theoretic model of F&J’s reference game, illustrated in Figure 1. The iterated best response (IBR) model of Franke (2009, 2011) applied to this game makes straightforward categorical predictions about the optimality of the shapes given a one-word message. The set-up for the signaling game associated with this task is given in Figure 2. The speaker, \mathcal{S} , has a type, either GS (the speaker wants to convey ‘green square’), GC (‘green circle’) or BC (‘blue circle’). \mathcal{S} sends a one-word message m (where m_G stands for “green”, and so on) to the hearer, \mathcal{H} , who must guess \mathcal{S} ’s type. If the guessing is successful, utility for both players is 1, else 0. An equilibrium is reached by first assuming a literal hearer \mathcal{H}_0 who only assumes that $\llbracket m \rrbracket$ denotes a true property of the object associated with her type (as denoted by the check marks in Figure 2). The speaker \mathcal{S}_1 considers which message(s) maximize(s) the probability of \mathcal{H}_0 guessing correctly. Then, \mathcal{H}_2 considers what \mathcal{S}_1 would do to determine the best response to each possible message. This reasoning iterates until a stable set of best responses has been reached. Below is a derivation showing that by the third iteration, we reach a unique set of best responses for each message: The hearer should respond to “blue” by guessing the blue circle, to “square” by guessing the green square, and to either “green” or “circle” by guessing the green circle. Intuitively, the green circle is the optimal choice given those messages because the hearer knows the speaker knows the hearer knows, etc., that if the speaker had wanted to convey ‘green square’, she’d have said “square”, and similarly for ‘blue circle’ and ‘blue’.

- (1) $\mathcal{H}_0 \quad m_G \rightarrow GS \text{ or } GC; m_S \rightarrow GS; m_B \rightarrow BC; m_C \rightarrow GC \text{ or } BC$
 $\mathcal{S}_1 \quad GS \rightarrow m_S; GC \rightarrow m_G \text{ or } m_C; BC \rightarrow m_B$
 $\mathcal{H}_2 \quad m_G \rightarrow GC; m_S \rightarrow GS; m_B \rightarrow BC; m_C \rightarrow GC$

When it comes to its predictions about interpretation, this instantiation of IBR is both categorical and normative: For each message, there is a single interpretation that the hearer *should* make, i.e., a single interpretation that any perfectly optimal reasoner would arrive at. If the hearer hears either “green” or “circle”, she should arrive at the optimal interpretation ‘green circle’. Of course, we know that actual human beings are not perfectly optimal reasoners. This has been discussed in game theory under the umbrella of *bounded rationality* going back to Simon (1957). The question is not whether our decision making deviates from perfect rationality, but rather why and how. This work and other work applying rational speech act (RSA) models (Frank and Goodman 2012; Goodman and Stuhlmüller 2013) tackles this question by positing cognitively

oriented Bayesian models to account for gradient preferences for strategically optimal actions in pragmatics experiments. Specifically in the context of F&J’s reference game, we need to explain why listeners’ guesses about speaker meaning are probabilistically skewed toward the optimal candidate (the green circle) in the “green” condition, while still showing loyalty to the most perceptually salient candidate (the blue circle) in the “circle” condition. (We focus on interpretation for now, turning to the issue of message selection at the end of Section 3.) Consider two possible explanations.

1. As F&J (and the RSA approach more generally) suggest, the cognitive task itself is inherently probabilistic, and the probabilities of choosing one solution vs. another are weighted by both perceptual salience (encoded in the priors of a Bayesian model) and by strategic considerations (encoded in the likelihood).
2. Alternatively, subjects may exhibit, both at the group level and at the individual level, distinct processes that conspire to give the overall effect of probabilistic weighting.

Considering the alternative, first note that it is in principle possible that each subject purely employs either a non-optimal, salience-based decision-making heuristic, or else an optimal strategy that can be modeled with simple IBR, and that the mixture of the two produces a gradient of responses only at the group level. This is an empirical question, and the issue of between-subject variation in pragmatics experiments is raised by Franke and Degen (2015), who conclude that some such variation does indeed come into play. But I’d like to abstract away from that, and consider a second possibility: that even within the individual, there are distinct “winner-takes-all” decision-making heuristics that can in principle be “turned on” or “turned off”, which need not reflect any inherently probabilistic reasoning. For example, as shown in the next section, it is conceivable that subjects use pure visual salience to “zero in” on some inherently preferred referent prior to receiving or considering the speaker’s message. Then, only if that preferred referent is inconsistent with the message, would the subject “turn on” her strategic reasoning and consider the pragmatically optimal referent (as per IBR).

One may ask what the difference is, practically speaking. Why worry about whether to model a set of experimental results with a single Bayesian reasoner, or with a two-tiered reasoner that uses interacting “winner-takes-all” heuristics? After all, if we are restricting ourselves to Marr’s computational level of analysis (i.e., characterizing the nature of the problem to be solved, with no detailed claims about how that problem is solved in real time), and if the overall predictions about rate of guessing a given referent are the same, then we won’t be able to distinguish the two models empirically. Firstly, the predictions might not be the same when applied to different sets of experimental results. Secondly, if a multi-tiered approach were to prove empirically successful, it would have the advantage of providing a convenient starting point for an algorithmic-level investigation, in which models could be tested against real-time behavioral data obtained via eye tracking or other methods in order to probe the actual cognitive mechanisms involved.

Apart from all that, there is a philosophical reason to at least consider a framework based on distinct interacting heuristics for modeling gradiently rational pragmatic behavior. In a recent blog post discussing gradience of grammaticality judgments, Hornstein (2015) draws an analogy: To pursue categorical models of grammaticality, like Minimalism (Chomsky 1995), in the face of gradient grammaticality judgments in the lab is merely to recognize that speakers’ judgments are the result of a complex decision-making process which can be decomposed into multiple components (one of which is a grammar), just as in physics a resultant force may be decomposed into a system of forces acting in tandem to produce a single effect. In the following section I outline a simple toy model of F&J’s toy data to show what it would look like, in principle, to treat interpretation behavior in the reference game as a resultant force and try to break it down into its component parts.

2 Salience first, ask questions later

F&J emphasize in the discussion of their reference game results that the modeling is only meant to be illustrative—this simple experiment alone is not expected to uncover any magic bullet to completely explain the relevant phenomena. I’d like to echo that sentiment; the model presented here is somewhat whimsical, and not meant to prove or disprove any particular claims. Rather, I want to sketch what the problem might look like from a different perspective, one where we do not assume that the task of the pragmatic reasoner is inherently a probability-assigning task.

To start, let’s construct a simple model of how one might go about choosing a shape in the “prior elicitation” condition, in which subjects receive no message and thus can choose a shape to their heart’s desire. F&J suggest, quite reasonably, that the probability of selection in this condition is directly determined by a gradient notion of salience. But we might be able to do away with these numerical values, and instead assume only a ranking of visual salience: The blue circle is more salient than the green square, and the green square is more salient than the green circle. (This is certainly in line with my own intuitions about which shapes “pop out” the most.) After that we just need a simple two-step procedure. First, select some random non-empty subset of the visual world (1–3 shapes) to give prior attention to. Then, select the most salient member of that subset as the “favored shape”, and simply pick that one. In other words, this model quickly narrows attention down to some part of the visual field, then zeroes in on the most visually salient or interesting shape in that part of the field. The ranking tells us that if the field contains the blue circle, the blue circle always wins. If the subject is selecting from among the two green shapes, then the green square wins. Only when the subject immediately attends only to the green circle does that one become the favorite. This is spelled out in the following algorithm. Note that though we present these models in the form of a simple algorithm, we are not conflating computational and algorithmic levels of analysis. We do not suggest that these models reflect the exact mechanisms at play. Rather, we suggest that this might provide a useful approximation of a purely salience and attention-based selection process which is active here.

- (2) Let \mathcal{V} denote the visual world, i.e., the set containing ‘green square’ (GS), ‘green circle’ (GC) and ‘blue circle’ (BC), and let $S(\mathcal{V})$ be the set of possible non-empty subsets of \mathcal{V} , i.e., $\mathcal{P}(\mathcal{V}) \setminus \{\}$.
 1. Randomly select a set of attended objects V from $S(\mathcal{V})$.
 2. Let v^* be the most salient object in V , where $\text{salience}(BC) > \text{salience}(GS) > \text{salience}(GC)$.
 3. Output v^*

This model works quite nicely for this data set in that it generates close approximations of the prior values reported by F&J—rather than assuming those values—without any requiring any numerical parameters. Where F&J report elicited priors of 0.3, 0.12, and 0.58 for GS , GC and BC , respectively, our parameter-free selection procedure predicts values of 0.29, 0.14 and 0.57. We only need to assume a relative ranking of visual salience.

Now we move on to the message conditions, where subjects were asked to interpret either “green” or “circle” as referring to one of the shapes. We start with the idea that hearers default to a method of selecting favorites as in (2). Then, they must assess whether their favored shape is consistent with the semantic content of the message they received from the speaker. If it is, i.e., if given the message m , $\llbracket m \rrbracket$ is a true property of their favored referent, then they simply choose that referent. For example, if the subject is already attending to the blue circle, and the message is “circle”, then the subject will simply guess ‘blue circle’, since the message is consistent with what they are already biased toward. But what if the message is “green” in that case? Then, under this model, the subject “zooms out” and assesses which of the three objects is most likely to have been referred to by a pragmatically competent hearer. This can be categorical, modeled as the selection of the equilibrium best response to the received message m in the IBR game in Figure 2.

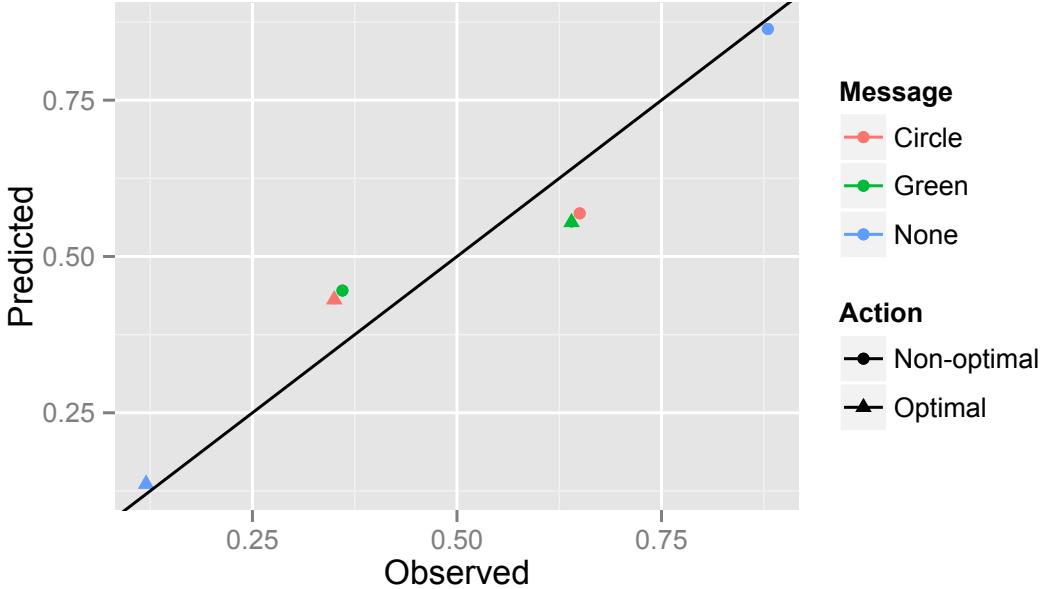


Figure 3: Prediction of observed interpretation data of a two-tiered model involving categorical optimality. Defection is not possible. (“Optimal” in this case refers to guesses of the referent that is optimal under an IBR model given a “green” or “circle” message, which is in both cases the referent ‘green circle’; “Non-optimal”, then, is the sum of guesses of the green square and the blue circle.)

- (3) Let \mathcal{V} denote the visual world, and let $S(\mathcal{V})$ be the set of possible non-empty subsets of \mathcal{V} .
 1. Randomly select a set of attended objects V from $S(\mathcal{V})$.
 2. Let v^* be the most salient object in V , where $\text{salience}(BC) > \text{salience}(GS) > \text{salience}(GC)$.
 3. If $\llbracket m \rrbracket(v^*) = \text{TRUE}$, set preferred interpretation i^* equal to v^* ; else, set i^* equal to $EBR(m)$, where EBR is the equilibrium best response in the game given in Figure 2.
 4. Output i^* .

At this point, we are still parameter-free, and we have not allowed for any noisy variation in the model. Therefore, we wouldn’t expect the predicted to mirror the observed values exactly. And indeed they don’t. However, as we see in Figure 3, the predictions are not wildly far off, either.

With a single adjustment, including the addition of a single free parameter, the predicted values line up almost perfectly with observation. To allow for greater variation, we add a step where the listener, if they have chosen the equilibrium best response as their favored referent, will *defect* from that strategy with a positive probability ϵ . Defection in this case means repeating the same salience-based strategy as before, i.e., attending again to some subset of the visual world, picking a favorite based on salience ranking, and then, if that favorite is consistent with m , guessing it.

- (4) Let \mathcal{V} denote the visual world, and let $S(\mathcal{V})$ be the set of possible non-empty subsets of \mathcal{V} .
 1. Randomly select a set of attended objects V from $S(\mathcal{V})$.

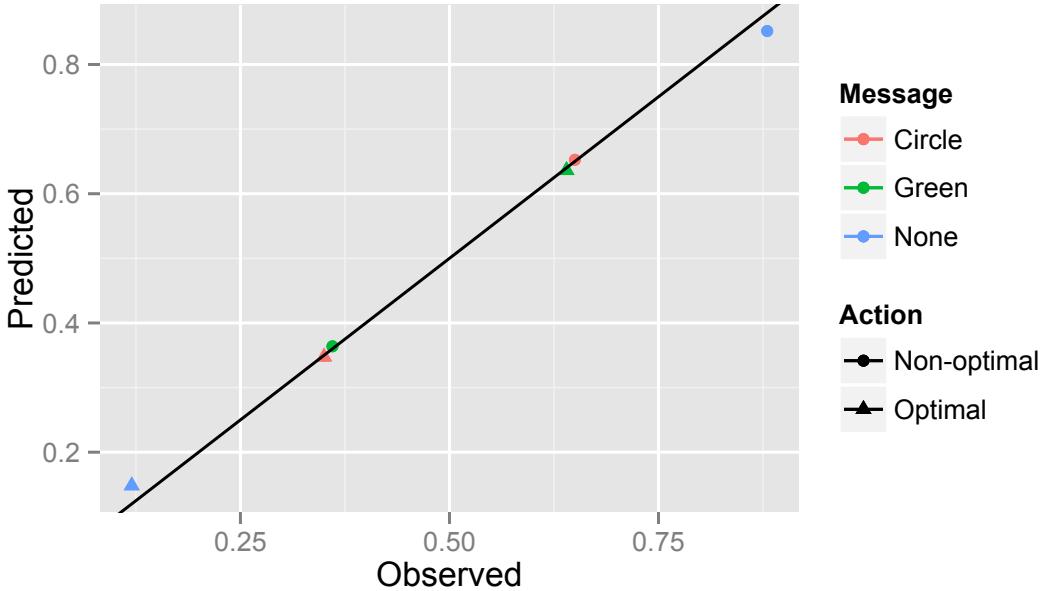


Figure 4: A two-tiered model with probability $\epsilon = 0.52$ of defecting from the optimal interpretation back to a purely salience-based interpretation strategy.

2. Let v^* be the most salient object in V .
3. If $\llbracket m \rrbracket(v^*) = \text{TRUE}$, set i^* equal to v^* ; else, set i^* equal to $EBR(m)$.
4. With probability ϵ , defect from $EBR(m)$: Select another set of attended objects V' , let $v^{*\prime}$ be the most salient member of V' , and if $\llbracket m \rrbracket(v^{*\prime}) = \text{TRUE}$, set i^* equal to $v^{*\prime}$
5. Output i^* .

This single-parameter model can be made to almost exactly reflect the reported observed data (see Figure 4) if we set ϵ to be quite high, at 0.52. That is, under this model, about half of subjects will respond to an inconsistency between message and favored referent by simply trying again. This paints a picture of a hearer who, at least in the context of this simple low-stakes Mechanical Turk experiment, puts off having to think strategically in favor of just picking something that works. Nonetheless, it is clear that some notion of strategic optimality is required to account for the data. I would not argue against an approach to pragmatic modeling based on strategic reasoning, nor against an approach based on Bayesian reasoning per se. Even the “winner-takes-all” maximization of expected utility in a signaling game will necessarily invoke Bayes’ rule in order to calculate the expected utility values, and although the Bayesian calculation may be simplified beyond recognition, as is the case with the IBR model invoked here, the Bayesian nature of the expected utility function is more obvious and important for signaling models of more complex phenomena (see Stevens et al. 2015, for one example). Here, we are merely taking seriously the question of whether expected-utility maximizing considerations might, for certain tasks, come into play only as a fallback, and what that would mean for a cognitively oriented approach to modeling these tasks.

So far, we have considered only the hearer’s interpretation of the speaker’s message. How does strategy

come into play during the selection of that message on the part of the speaker? F&J’s observed values are quite a bit closer to what we’d expect from our simple IBR model, which in this case derives a categorical best message for the green square (“square”) and the blue circle (“blue”), and predicts random selection between the two messages for the green circle. For the green shapes, the production data are within 0.06 of the predicted values, but for the blue circle, the message “circle” is chosen at a higher-than-expected rate of 0.17. Taking the same approach as we did above for interpretation, we can try a simple two-tiered procedure incorporating visual salience. First, as before, choose a subset of the visual world $S(\mathcal{V})$, then focus attention on the most salient member of that subset, the “favored” shape. If the favored shape happens to be the shape for which the speaker is instructed to provide a one-word label, then the speaker simply gives a single property that distinguishes the favored shape from any other shapes in $S(\mathcal{V})$. For example, if $S(\mathcal{V})$ is {‘green square’, ‘blue circle’}, then the speaker could choose either “blue” or “circle” to label the blue circle, because either of those properties would distinguish the intended referent from the other member of $S(\mathcal{V})$. And just as before, if a non-favored shape is assigned, the speaker zooms out and considers which label is strategically optimal as per an IBR model. Even without any defection, the output of this model comes close to the observed values: In the ‘green square’ condition, the rate of selecting “square” is predicted to be 0.93, vs. 0.07 for “green”, in the ‘green circle’ condition, the model predicts random selection of either “green” or “circle”, and in the ‘blue circle’ condition, the rate of selection “blue” is predicted to be 0.85 vs. 0.15 for “circle”. This doesn’t explain the bias toward “circle” over “green” in the ‘green circle’ condition, but I am unconvinced that F&J’s inclusion of a shape-word bias in their model, which without independent justification seems purely descriptive, offers any better explanation. If we allowed ourselves the addition of a similar parameter, we would surely get a closer fit.

I end this section by pointing out that the nature of the experimental task itself plays a crucial role in modeling observations obtained in the laboratory. I have suggested that for this particular experiment, the amount of strategy being played, as opposed to dumber forms of reasoning, is perhaps quite low. But that is surely in part due to the forced-choice nature of the task, which we should expect to favor decision-making mechanisms that check prior biases against new information and stick with those biases as long as they remain consistent. This contrasts with Frank & Goodman (2012), where subjects are asked to make monetary bets on what was meant by a message, or on what word a speaker will use to describe a particular object. For a task based on betting, subjects are essentially told to reason explicitly about probabilities. Thus we shouldn’t expect the behavior in that reference game experiment to mirror the behavior in F&J’s. For example, we might expect subjects in Frank and Goodman’s experiment to hedge their bets, resulting in less overall difference between the shapes/messages. And indeed, the Frank and Goodman’s Figure 1 shows subjects’ bets to be clustered around the safe bets of 50 and 33 (of 100). Further experimentation could probe whether certain task-specific differences in predictions are fully borne out.

3 When do we think strategically?

One of the hallmarks of the rational speech act approach to pragmatic analysis is a move from purely normative models of pragmatic optimality (like most game-theoretic accounts) toward cognitive models of observed pragmatic behavior. This is a valuable endeavor, but by dumping a number of distinct psychological variables, notions of salience, attention, cognitive biases, etc., into a single Bayesian formula for each speaker/hearer type, one wonders whether the approach stops short of its goal of providing cognitive reasons for observed phenomena, instead occupying a middle ground where the relevant phenomena are simply described. While concise description can contribute to explanation, the explanation itself should attempt to suss out whether the described phenomenon emerges from the interaction of distinct behaviors. This desideratum

is independent of Marr levels: Whether we are at the computational level (or in F&J's hierarchy, the level of reasons) or the algorithmic level (the level of processes), we should strive to achieve an appropriate level of granularity in our explanations. To that end, we must ask the question of whether strategic and non-strategic thinking both play independent roles in explaining pragmatic behavior. And to the extent that the two can be teased apart, we are likely to see that the most elegant models of the strategic side of this dichotomy are the good old-fashioned game-theoretic models which have already enjoyed a fair amount of attention in the last decade (Benz, Jäger, and van Rooij 2006; Franke 2009; Clark 2011).

That brings us to our final question: When do we think strategically? It might be useful to first consider what it means to think strategically. Let's do so in the context of F&J's carpet negotiation scenario, their third test case for probabilistic pragmatics. In this scenario, a prospective buyer of a carpet has either moderate, strong or no interest in owning the carpet, and the seller has either moderate or strong interest in selling it. The buyer can signal her interest to buy directly ("I want to buy this carpet") or be indirect ("this is a beautiful carpet") in order to get a lower price offer from the seller. F&J model this scenario game-theoretically, and conclude that the indirect strategy is optimal "if [the speaker's] stakes are low or if she believes that her opponent's stakes are high." (p.24). That is, it is not always a smart move to simply declare, "I want to buy this carpet", as the buyer could get a lower price by being more coy. For this section, F&J return to normative models: It is clear that theirs is a model of how a carpet negotiation *should* go, and that we would expect to find incompetent price negotiators out in the real world. With this in mind, imagine that a prospective buyer does in fact employ the strategically optimal utterance, and call the thinking that leads to this utterance "strategic thinking". This strategic thinking could take two forms.

1. Consciously strategic: The buyer ascertains that the seller is eager to sell, and reasons that if she indirectly signals mild interest in the carpet, the seller might undervalue her interest and offer a favorable price in order to sell the carpet faster.
2. Unconsciously strategic: The buyer grew up in a family of expert negotiators and learned by observation that whenever a seller appeared eager to sell, her family members always opened price negotiations by saying something vaguely nice about the item such as, "this is a beautiful carpet."

In the first case, the normative reasoning model is explicitly represented in the mind of our buyer. In the second case, the buyer might not have any thoughts about why she is doing what she is doing, yet the behavior is still strategic in that it has its roots in a form of optimality: The success of certain negotiation strategies led to the development of norms and conventions in the buyer's community which could be learned from observation.

Based on this broad categorization, we can re-frame our question as two questions: When do people (fail to) apply strategies they've learned, and when do people (fail to) work out the best strategy for themselves? It's easy to imagine strategic thinking failing at both levels. We could imagine even an experienced buyer being overcome with delight upon seeing a carpet and gleefully exclaiming, "I want to buy this carpet!" Should such an act be seen as the result of inherently probabilistic reasoning which balances strategic considerations with psychological biases? Surely not. Rather, notions like expected utility simply don't come into play in that case. It is along the same lines that I raise the possibility that a forced-choice reference game will give rise to reasoning behaviors either partially or totally devoid of strategic considerations. But in any case, we know that pragmatic behavior does seem to skew toward strategic optimality, and thus normative models can still wield some explanatory power in cognitive science, even if they only paint a partial picture.

References

- Benz, Anton, Gerhard Jäger & Robert van Rooij (eds.). 2006. *Game theory and pragmatics*. New York: Palgrave MacMillan.
- Chomsky, Noam. 1995. *The minimalist program*. Cambridge, MA: MIT Press.
- Clark, Robin. 2011. *Meaningful games*. Cambridge, MA: MIT Press.
- Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998. doi:10.1126/science.1218633.
- Franke, Michael. 2009. *Signal to act: Game theory in pragmatics*. Amsterdam: Universiteit van Amsterdam Dissertation.
- Franke, Michael. 2011. Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantic and Pragmatics* 4(1). 1–80. doi:10.3765/sp.4.1.
- Franke, Michael & Judith Degen. 2015. Reasoning in reference games: Individual- vs. population-level probabilistic modeling. Ms., Tübingen & Stanford.
- Franke, Michael & Gerhard Jäger. This volume. Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft* 35(1). xx–xx.
- Goodman, Noah D & Andreas Stuhlmüller. 2013. Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science* 5(1). 173–184.
- Hornstein, Norbert. 2015. Degrees of grammaticality? <http://facultyoflanguage.blogspot.com/2015/09/degrees-of-grammaticality.html>.
- Marr, David. 1982. *Vision: A computational approach*. San Francisco, CA: Freeman & Co.
- Parikh, Prashant. 2010. *Language and equilibrium*. Cambridge, MA: MIT Press.
- Simon, Herbert A. 1957. *Models of man*. New York: Wiley.
- Stevens, Jon, Anton Benz, Sebastian Reuß & Ralf Klabunde. 2015. A strategic reasoning model for generating alternative answers. In Chenggong Zong & Michael Strube (eds.), *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Vol. 1: Long Papers)*, 534–542. Beijing, China: Association for Computational Linguistics. <http://www.aclweb.org/anthology/P15-1052>.