The iterated best response model for referential games

Chris Potts, Ling 236: Context dependence in language and communication, Spring 2012

May 21

Overview

Quick history

- Lewis (1969) is the earliest decision-theoretic approach to pragmatics. Lewis was interested primarily in how linguistic conventions arise in a speech community. To pursue this question, he developed signaling games. See also Lewis 1975.
- Lewis's work was largely ignored in linguistics for decades. The sole exception is Clark (1996) and related work, which anticipates many of the developments to come.
- In the 1990s, Prashant Parikh proposed a game-theoretic model for linguistic interactions (semantics and pragmatics). See especially Parikh 2000, 2001.
- Blutner (1998, 2000) developed Bidirectional Optimality Theory, which formalized and extended insights by Horn (1984) about how balancing speaker effort and hearer effort can give to stable kinds of pragmatic enrichment. Jäger (2002) is an elegant formal restatement of the theory with a number of extensions building on insights by Frank and Satta (1998) and Karttunen (1998).
- van Rooy (2003, 2004) sought to reformulate and extend Blutner ideas using Lewisian signaling system. Those papers point out some fundamental limitations of Bi-OT and argue that signaling games offer a more general solution to problems related to pragmatic enrichment.
- Signaling systems have too many equilibria, not all of them intuitively alike. The Iterated Best Response models of Franke (2008, 2009) and Jäger (2007, To appear) seek to address this.

This handout

- This handout builds on the presentation of signaling games in Jäger (To appear), though I made various simplifications in order to keep the presentation manageable. (Basically, my games leave out the cost functions on forms and do not introduce epistemic indeterminacy.)
- My primary goal is to see how these games work for various referential tasks.
- In addition to all the models listed above, we also have Golland et al. (2010) and Frank and Goodman (2012). Very little is known right now about how all these models relate to each other. I think the best we can is to try to find places where they agree or disagree.

1 The model

1.1 Signaling games

Definition 1 (Pure communication cheap-talk signaling systems). A pure communication cheap talk signaling system is a tuple $\langle s, h, \Omega, M, [\![\cdot]\!], P, \sigma, \rho, u\rangle$, where

i. s is the speaker, and h is the hearer
ii. Ω is a set of states
iii. M is a set of messages
iv. $[\![\cdot]\!] : S \mapsto \wp(\Omega)$ is the semantic interpretation function
v. P is prior probability distribution over Ω
vi. σ is a function from worlds into probability distributions over signals: h's expectations about s
vii. ρ is a function from signals into probability distributions over worlds: s's expectations about h
viii. $u : (\Omega \times \Omega) \mapsto \mathbb{R}$ is a utility function defined so that $u(w_i, w_j) = 1$ if $w_i = w_j$, else 0.

For the following definitions, assume a single signaling system $\mathcal{G} = \langle s, h, \Omega, M, [\![\cdot]\!], P, \sigma, \rho, u\rangle$.

Definition 2 (Speaker expected utility). s's expected utility for $m \in M$ given $w \in \Omega$ is defined in terms of the utilities and how the *receiver* acts given m:

$$EU_s(m \mid w) \stackrel{def}{=} \sum_{w' \in \Omega} \rho(w'|m) * u(w, w')$$

Definition 3 (Speaker best response to a world). s's best response when presented with a world w is the set of all signals that maximize expected utility given w:

$$br_s(w) \stackrel{def}{=} \{m \mid EU_s(m \mid w) = \max_{m' \in M} EU_s(m' \mid w)\}$$

Definition 4 (Speaker posterior). The posterior for σ is defined via Bayes' rule:

$$\sigma(w \mid m) \stackrel{def}{=} \frac{\sigma(m|w) * P(w)}{\sum_{w' \in \Omega} \sigma(m|w') * P(w')} \quad \text{undefined where the denominator is 0}$$

Definition 5 (Hearer expected utility). h's expected utility for $w \in \Omega$ given $m \in M$ is defined in terms of the utilities and how the *speaker* acts given w:

$$EU_h(w \mid m) \stackrel{def}{=} \begin{cases} \sum_{w' \in \Omega} \sigma(w' \mid m) * u(w, w') & \text{if defined} \\ \text{undefined} & \text{otherwise} \end{cases}$$

Definition 6 (Hearer best responses). h's best response when presented with a signal m is the set of all worlds that maximize expected utility given m, resorting to the literal interpretation of m where the signal conflicts with σ:

$$br_h(m) \stackrel{def}{=} \begin{cases} \{w \mid EU_h(w \mid m) = \max_{w' \in \Omega} EU_h(w' \mid m)\} & \text{if } EU_h(\cdot \mid m) \text{ is defined} \\ [\![m]\!] & \text{otherwise} \end{cases}$$

1.2 Iterated best responses

Definition 7 (Speaker response strategy). s's best response strategy to h's strategy ρ is the one that assigns equal probability to all best responses:

$$BR_s(\rho) = \sigma \text{ iff } \sigma(m \mid w) = \frac{|\{m\} \cap br_s(w)|}{|br_s(w)|} \text{ for all } w \in \Omega \text{ and } m \in M$$

Definition 8 (Hearer response strategy). h's best response to s's strategy σ is the one that conditions on the best responses:

$$BR_h(\sigma) = \rho \text{ iff } \rho(w \mid m) = \frac{P(\{w\} \cap br_h(m))}{P(br_h(m))} \text{ for all } w \in \Omega \text{ and } m \in M$$

2 Examples

2.1 Simple scalar implicature (adapted from Jäger To appear)

(1)
a. $\Omega = \{w_{\neg\exists}, w_{\exists\neg\forall}, w_{\forall}\}$

b. $M = \{$"none", "some", "all"$\}$

c. $[\![$"none"$]\!] = \{w_{\neg\exists}\}$; $[\![$"some"$]\!] = \{w_{\exists\neg\forall}, w_{\forall}\}$; $[\![$"all"$]\!] = \{w_{\forall}\}$

d. P = even distribution over worlds

e. Suppose the initial speaker strategy is to be honest and literal:

$\sigma =$

	"none"	"some"	"all"
$w_{\neg\exists}$	1	0	0
$w_{\exists\neg\forall}$	0	1	0
$w_{\forall}$	0	½	½

f. Then $BR_h(\sigma)$ uses the largest value from each column:

$BR_h(\sigma) =$

	$w_{\neg\exists}$	$w_{\exists\neg\forall}$	$w_{\forall}$
"none"	1	0	0
"some"	0	1	0
"all"	0	0	1

g. The speaker's best response follows the same principle:

$BR_s(BR_h(\sigma)) =$

	"none"	"some"	"all"
$w_{\neg\exists}$	1	0	0
$w_{\exists\neg\forall}$	0	1	0
$w_{\forall}$	0	0	1

h. The strategies have now stabilized.

2.2 Golland et al. (2010)

The core model of Golland et al. is a signaling system. Their p is the prior P. Their p_S is the sender strategy. Their p_L is the hearer strategy. Their U is the utility function (defined as in our pure communication games). Their embedded model is just a round of iterated best responses. The only difference, as far as I can tell, is that they define a single expected utility function incorporating both speaker and hearer strategies. I am not sure what consequences this has.

(2) a. $\Omega = \{r_{\text{vase}}, r_{\text{table}}\}$

b. $M = \{\text{'right of lamp'}, \text{'on table'}\}$

c. $[\![\text{'right of lamp'}]\!] = \{r_{\text{vase}}, r_{\text{table}}\}$; $[\![\text{'on table'}]\!] = \{r_{\text{vase}}\}$

d. P is an even distribution over the referents in Ω

e. $\sigma =$

	'on table'	'right of lamp'
r_{table}	0	1
r_{vase}	½	½

f. $BR_h(\sigma) =$

	r_{table}	r_{vase}
'on table'	0	1
'right of lamp'	1	0

g. $BR_s(BR_h(\sigma)) =$

	'on table'	'right of lamp'
r_{table}	0	1
r_{vase}	1	0

h. The strategies have now stabilized.

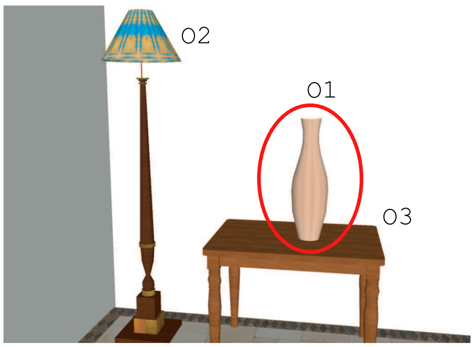

Figure 1: An example of a 3D model of a room. The *speaker*'s goal is to reference the target object O1 by describing its spatial relationship to other object(s). The *listener*'s goal is to guess the object given the speaker's description.

2.3 Frank and Goodman (2012)

(3) a. $\Omega = \{r_{\text{blue square}}, r_{\text{blue circle}}, r_{\text{green square}}\}$

b. $M = \{$“blue”, “square”, “circle”, “green”$\}$

c. $[\![$“blue”$]\!] = \{r_{\text{blue square}}, r_{\text{blue circle}}\}$ $[\![$“square”$]\!] = \{r_{\text{blue square}}, r_{\text{green square}}\}$
$[\![$“circle”$]\!] = \{r_{\text{blue circle}}\}$ $[\![$“green”$]\!] = \{r_{\text{green square}}\}$

d. P is an even distribution over the referents in Ω (changing P seems to have no effect).

e. $\sigma =$

	“blue”	“square”	“circle”	“green”
$r_{\text{blue square}}$	½	½	0	0
$r_{\text{blue circle}}$	½	0	½	0
$r_{\text{green square}}$	0	½	0	½

f. $BR_h(\sigma) =$

	$r_{\text{blue square}}$	$r_{\text{blue circle}}$	$r_{\text{green square}}$
“blue”	½	½	0
“square”	½	0	½
“circle”	0	1	0
“green”	0	0	1

g. $BR_s(BR_h(\sigma)) =$

	“blue”	“square”	“circle”	“green”
$r_{\text{blue square}}$	½	½	0	0
$r_{\text{blue circle}}$	0	0	1	0
$r_{\text{green square}}$	0	0	0	1

h. $BR_h(BR_s(BR_h(\sigma))) =$

	$r_{\text{blue square}}$	$r_{\text{blue circle}}$	$r_{\text{green square}}$
“blue”	1	0	0
“square”	1	0	0
“circle”	0	1	0
“green”	0	0	1

i. The strategies have now stabilized.

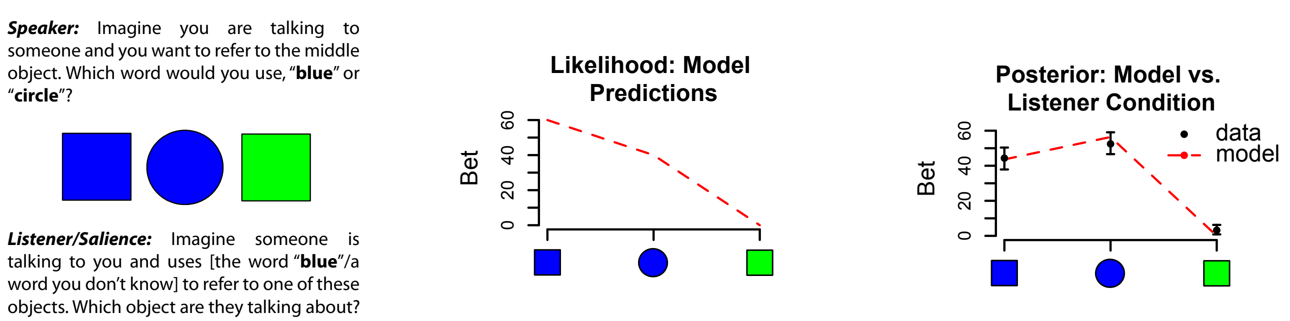

Conclusion The experimental results are consistent with the claim that people differ with regard to whether they bother to reason to the stable strategy. The priors don’t play a role. (For discussion of why this might be, see Franke 2009:§3.1.)

2.4 Stiller et al. (2011)

The scalar condition of Stiller et al. (2011) is isomorphic to that of Frank and Goodman (2012). It's worth looking at their 'no scales' condition, though. Stiller et al. report that children and adults are not able to lean the relevant associations without strong priors. The IBR model captures the relevant assocations.

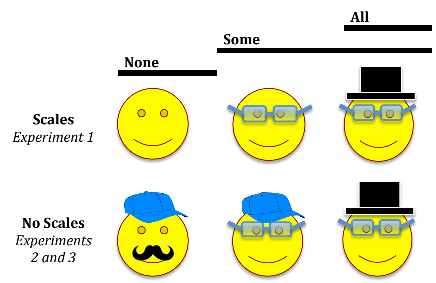

(4) a. $\Omega = \{r_{\text{cap+mustache}}, r_{\text{cap+glasses}}, r_{\text{tophat+glasses}}\}$

b. $M = \{$"cap", "glasses", "mustache", "tophat"$\}$

c. $[\![$"cap"$]\!] = \{r_{\text{cap+mustache}}, r_{\text{cap+glasses}}\}$ $[\![$"glasses"$]\!] = \{r_{\text{cap+glasses}}, r_{\text{tophat+glasses}}\}$
$[\![$"mustache"$]\!] = \{r_{\text{cap+mustache}}\}$ $[\![$"tophat"$]\!] = \{r_{\text{tophat+glasses}}\}$

d. P is an even distribution over the referents in Ω (changing P seems to have no effect).

e. $\sigma =$

	"cap"	"glasses"	"mustache"	"tophat"
$r_{\text{cap+mustache}}$	½	0	½	0
$r_{\text{cap+glasses}}$	½	½	0	0
$r_{\text{tophat+glasses}}$	0	½	0	½

f. $BR_h(\sigma) =$

	$r_{\text{cap+mustache}}$	$r_{\text{cap+glasses}}$	$r_{\text{tophat+glasses}}$
"cap"	½	½	0
"glasses"	0	½	½
"mustache"	1	0	0
"tophat"	0	0	1

g. $BR_s(BR_h(\sigma)) =$

	"cap"	"glasses"	"mustache"	"tophat"
$r_{\text{cap+mustache}}$	0	0	1	0
$r_{\text{cap+glasses}}$	½	½	0	0
$r_{\text{tophat+glasses}}$	0	0	0	1

h. $BR_h(BR_s(BR_h(\sigma))) =$

	$r_{\text{cap+mustache}}$	$r_{\text{cap+glasses}}$	$r_{\text{tophat+glasses}}$
"cap"	0	1	0
"glasses"	0	1	0
"mustache"	1	0	0
"tophat"	0	0	1

i. The strategies have now stabilized.

2.5 Deeper?

- The implicature example stabilized after 1 hearer strategy and 2 speaker strategies.
- The Golland et al. example stabilized after 1 hearer strategy and 2 speaker strategies.
- The Frank and Goodman example stabilized after 2 hearer strategies and 2 speaker strategies.
- Strikingly, scenarios like the following stabilize at the same rate as Frank and Goodman's:

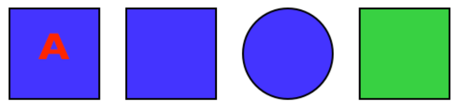

i. $BR_s(BR_h(\sigma)) =$

	"blue"	"square"	"circle"	"green"	"contains A"
$r_{\text{blue square with A}}$	0	0	0	0	1
$r_{\text{blue square}}$	½	½	0	0	0
$r_{\text{blue circle}}$	0	0	1	0	0
$r_{\text{green square}}$	0	0	0	1	0

ii. $BR_h(BR_s(BR_h(\sigma))) =$

	$r_{\text{blue square with A}}$	$r_{\text{blue square}}$	$r_{\text{blue circle}}$	$r_{\text{green square}}$
"blue"	0	1	0	0
"square"	0	1	0	0
"circle"	0	0	1	0
"green"	0	0	0	1
"contains A"	1	0	0	0

- Can we think of referential games that require even deeper reasoning? The (im)possibility of such examples might help decide between Frank and Goodman's model and the IBR model.

3 Other remarks

- σ is the hearer's model of the speaker, and ρ is the speaker's model of the hearer. This cross-over is reminiscent of the way interrogator and expert utilities are intermingled by Benz (2005).
- Franke (2009) and Jäger (To appear) explore models in which messages have costs. This makes it possible to account for manner-based implicatures like those deriving from the principle that (ab)normal things are described with (ab)normal language.
- The above presentation delivers implicatures relentlessly. Franke (2009) shows how to weaken this by introducing epistemic uncertainly about which game is being played. See also Jäger (To appear).
- Jäger (To appear) argues that the IBR model is compatible with alternative generation in the manner of Chierchia et al. (To appear). One way to read this is that the IBR approach is independent of the grammatical approach. They can even be seen as complementary, with the grammatical approach providing logical forms that we can reason about game-theoretically.

References

Benz, Anton. 2005. Utility and relevance of answers. In Anton Benz; Gerhard Jäger; and Robert van Rooij, eds., *Game Theory and Pragmatics*, 195–219. Basingstoke, Hampshire: Palgrave McMillan.

Blutner, Reinhard. 1998. Lexical pragmatics. *Journal of Semantics* 15(2):115–162.

Blutner, Reinhard. 2000. Some aspects of optimality in natural language interpretation. *Journal of Semantics* 17(3):189–216.

Chierchia, Gennaro; Danny Fox; and Benjamin Spector. To appear. The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In Maienborn et al. (To appear).

Clark, Herbert H. 1996. *Using Language*. Cambridge: Cambridge University Press.

Frank, Michael C. and Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* .

Frank, Robert and Giorgio Satta. 1998. Optimality theory and the generative complexity of constraint violability. *Computational Linguistics* 24(2):307–315.

Franke, Michael. 2008. Interpretation of optimal signals. In Krzysztof R. Apt and Robert van Rooij, eds., *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logics and Games*, 297–310. Amsterdam University Press.

Franke, Michael. 2009. *Signal to Act: Game Theory in Pragmatics*. ILLC Dissertation Series. Institute for Logic, Language and Computation, Universiteit van Amsterdam.

Golland, Dave; Percy Liang; and Dan Klein. 2010. A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, 410–419. Cambridge, MA: Association for Computational Linguistics. URL `http://www.aclweb.org/anthology/D10-1040`.

Horn, Laurence R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In Deborah Schiffrin, ed., *Meaning, Form, and Use in Context: Linguistic Applications*, 11–42. Washington: Georgetown University Press.

Jäger, Gerhard. 2002. Some notes on the formal properties of bidirectional Optimality Theory. *Journal of Logic, Language and Information* 11(4):427–451.

Jäger, Gerhard. 2007. Game dynamics connects semantics and pragmatics. In Ahti-Veikko Pietarinen, ed., *Game Theory and Linguistic Meaning*, 89–102. Amsterdam: Elsevier.

Jäger, Gerhard. To appear. Game theory in semantics and pragmatics. In Maienborn et al. (To appear).

Karttunen, Lauri. 1998. The proper treatment of optimality in computational phonology. Technical Report, Xerox ResearchCentreEurope.

Lewis, David. 1969. *Convention*. Cambridge, MA: Harvard University Press. Reprinted 2002 by Blackwell.

Lewis, David. 1975. Languages and language. In Keith Gunderson, ed., *Minnesota Studies in the Philosophy of Science*, volume VII, 3–35. Minneapolis: University of Minnesota Press. Reprinted in Lewis 1983, 163–188. Page references are to the reprinting.

Lewis, David. 1983. *Philosophical Papers*, volume 1. New York: Oxford University Press.

Maienborn, Claudia; Klaus von Heusinger; and Paul Portner, eds. To appear. *Semantics: An International Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter.

Parikh, Prashant. 2000. Communication, meaning, and interpretation. *Linguistics and Philosophy* 23(2):185–212.

Parikh, Prashant. 2001. *The Use of Language*. Stanford, CA: CSLI.

van Rooy, Robert. 2003. Relevance and bidirectional OT. In Reinhard Blutner and Henk Zeevat, eds., *Pragmatics in Optimality Theory*, 173–210. Oxford: Palgrave Macmillan.

van Rooy, Robert. 2004. Signalling games select Horn strategies. *Linguistics and Philosophy* 27(4):493–527.

Stiller, Alex; Noah D. Goodman; and Michael C. Frank. 2011. Ad-hoc scalar implicature in adults and children. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Societys*.