

Indirect answers and cooperation: On Asher and Lascarides's 'Making the right commitments in dialogue'

Christopher Potts
UMass Amherst

Abstract This commentary argues that linguistic cooperation is essential even in discourse situations in which the nonlinguistic preferences of the participants are misaligned. The central examples involve indirect answers to direct questions. The analysis builds on the work of Asher and Lascarides, without, though retreating from the axioms of cooperativity as hastily as they do in the workshop paper (Asher & Lascarides 2008). I also argue (section 4) that discourse coherence and inferences from the common ground can account for much pragmatic enrichment.

1 The importance of coherence

For the semantically-inclined pragmaticist — the sort who wants to stick close to language even while moving to the richer problems of message enrichment in social contexts — reading Asher & Lascarides's (2003) *Logics of Conversation* (henceforth LoC) is inspiring. The overall message is that if you get the discourse coherence relations right, and reason seriously in terms of contextual information, then enriched pragmatic meanings will emerge. The details are a grind, but the basic idea is compelling enough to keep one's spirits up.

I've only begun my study of LoC, but I see fairly clearly that cooperativity is part of its foundation. When we walk back through the book's important high-level generalizations, we often encounter cooperativity along the way. The discourse participants might want very different things in life, but, when they exchange information in dialogue, they have an incentive to coordinate, not only on the meanings of their words and phrases, but also on how those linguistic pieces are woven into a discourse.

It's thus surprising to find that 'Making the right commitments in dialogue' (henceforth 'Commitments') comes out boldly against this cooperative foundation. The abstract opens by saying that accounts like LoC's "fail to say anything about conversations where a robust notion of cooperation is absent". The conversations A&L use to motivate this are extremely (even absurdly) confrontational — so much so that it can be hard to imagine what they would involve except unbridled aggression or unchecked duplicity.

With this commentary, I urge caution when moving away from the cooperative foundation. It is conceivable that A&L are right that most pragmatic theories say little or nothing about situations so uncooperative that little information can be exchanged. These might not be very interesting cases for pragmatic theory, though. I argue that LoC and ‘Commitments’ have a lot to say about situations in which preferences are misaligned but the players are still attempting to communicate.

Indirect answers to direct questions are the empirical foundation of this commentary. The next section introduces the central tool in LoC for modeling such answers: the coherence relation of Indirect Question–Answer Pairs and its associated axiom. I emphasize their connection to cooperativity as defined in LoC. Section 3 reviews a range of indirect answer types, seeking, in particular, to build a case that these relations, though rooted in cooperativity, are alive and important even in discourse situations in which preferences are misaligned. As a group, I think these cases support the LoC view of linguistic cooperation.

Section 4 then steps back to assess what else needs to be done in order to deliver rich pragmatic inferences. I argue, with A&L, that the coherence relations, working in conjunction with inferences from contextual information, tell much of the story, and I review some very simple computational simulations to bolster that conclusion. Finally, section 5 raises two tangentially related issues that I personally would benefit from discussing with A&L and the other workshop participants.

2 The Indirect Question–Answer Pair (IQAP) axiom

The Indirect Question–Answer Pair (IQAP) axiom, stated in (1), makes only a brief appearance in ‘Commitments’, but it is central to LoC. I argue that it is central to the concerns of ‘Commitments’ as well, though this is perhaps not immediately evident.

$$(1) \quad \text{IQAP: } (\lambda : ?(\alpha, \beta) \wedge \text{int}(\alpha)) > \lambda : \text{IQAP}(\alpha, \beta)$$

‘Generally, if an interrogative α is related by some coherence relation $?$ to β (in discourse λ), then that relation is the indirect question–answer pair relation.’

IQAP establishes a defeasible connection between interrogatives and the utterances that follow them in discourse: if you ask me a question α , and I respond with β , then you will expect $\text{IQAP}(\alpha, \beta)$ to hold. I should, in turn, plan my utterance with your expectations in mind. You are aware that I will plan with these expectations in mind. I am aware that you’re aware of this. And so forth.

The result is stability in discourse coherence, albeit with room for exceptions.

In LoC, IQAP is not quite axiomatic. ‘Commitments’ doesn’t provide the details of its derivation, but it hints at them when it says that IQAP and Q-Elab “follow from their [LoC’s] formalization of axioms of rationality and cooperativity” (p. 11). It is worth expanding on this. In LoC, the primitive statement is Cognitive IQAP (p. 404), which depends on the beliefs and inferential capabilities of the questioner. A&L derive the simpler form (1) from Cognitive IQAP and other axioms of their system. For present purposes, the step that is most important is the one where they call upon Cooperativity (p. 391), which is stated informally in (2).

(2) Cooperativity

For two agents A and B in a single discourse:

- a. Generally, if A intends φ , then B intends φ .
- b. Generally, if A cannot adopt B’s intentions, then A intends to make this known to B.

The derivation of IQAP from its cognitive counterpart is important because it frees IQAP, in some sense, from its entanglements with deep pragmatic notions. A&L write:

while validating these axioms is achieved through reasoning about cognitive states (see Asher & Lascarides (2003)), the premises of these axioms don’t talk about beliefs or intentions at all, and instead rest entirely on sentence mood. (p. 11)

I take this to mean the following: if we dig down deep inside of IQAP, we find rich cognitive notions connected with cooperation, in the sense of shared intentions and openness that (2) captures. IQAP inferences might seem to be only form-based, but they are still mediated by genuinely pragmatic notions. Thus, changing the pragmatics should change the IQAP inferences.

The next section assesses this idea. Before moving on, though, it is worth briefly describing what the IQAP relation imposes on the connection between the question α and the answer β . Here is a prose adaptation of the statement and description in LoC (p. 403):¹

- (3) $IQAP(\alpha, \beta)$ holds only if there is a true direct answer p to the question $[[\alpha]]$, and the questioner can infer p from $[[\beta]]$ in the utterance context.

¹ I use $[[\cdot]]$ informally to pick out the interpretations of expressions.


I think some of the exchanges below call into question the emphasis on truth, but the inferential relationship between the direct and indirect answers is important, as is the orientation to the *questioner's* capabilities.

3 Indirect answers and non-cognitive IQAP

My empirical focus is on question–answer pairs in which the answer is semantically (but not pragmatically) unrelated to the question. Such indirect answers are central to the LoC account of the pragmatics of answerhood. A slight variation on the following example opens their discussion of IQAP (p. 313):


- (4) a. A: Did John pass his exams?
b. B: He got 60%.

Let's suppose that A's question denotes the set $\{p, \neg p\}$, where p is the proposition that John passed his exams and $\neg p$ is its complement. In general logical terms, B does not fully address the issue raised by A: his answer, call it q , has a nonempty intersection with both p and $\neg p$, which means that q fails to resolve the issue (Groenendijk & Stokhof 1984; Roberts 1996; van Rooy 2004):²



A&L note that B seems not to have resolved A's question. They continue, though, by saying that A might well be able to reason in terms of what he knows in order to turn B's superficially partial answer into a resolving one. Suppose A knows that a passing grade is at or above 60%. Then B has answered A's question; in (6), I've simply wiped out the space in which $q \wedge \neg p$ is true, to reflect this world knowledge.

² It helps to imagine that we are simply unsure of the standard for failing. It could be 60%, but it could be higher or lower.



Learning q brings us into the gray space. Because the upper right corner has been eliminated, the gray space entails that John passed his exams. The necessary condition in (3) is satisfied, because the questioner can calculate that β , together with the common ground, entails the direct answer.

It is no surprise that answers like this can be resolving despite being superficially unrelated to the issue at hand. Speakers who use them are simply implicitly asking their listeners to reason in terms of other knowledge they have. It's a common discourse strategy. In the Switchboard example (7), Speaker A.101 relies on it as well.³

(7) B.100: So will, you, uh, breed one of the litter, then next year?

A.101: We're, we sold all that litter,

One can't breed what one has already sold. Thus, worlds in which A.101 has both sold the entire litter and breeds one of its members are outside the realm of possibility. As a result, Speaker B.100 likely felt that his question had been thoroughly answered.

Of course, indirect answers of this form are not simply direct answers in disguise. They are *over-answers*, i.e., they fully address the immediate question and they also raise and answer some additional issue. For example, a simple 'yes' answer to A's question in (4a) would take us to the left side of (6), but it would leave open which numerical grade John got. In effect, B's indirect answer resolves both A's question as well as the question of which numerical grade John got (as well as all its subquestions, e.g., Did John get 60%?; see [Büring 1999](#)).

If this were all that indirect answers involved, then there would be little need to invoke discourse coherence relations (or any of the Gricean stuff falling under the heading of Relevance). These are logical deductions; the questioners in (4) and (7) might as well have looked their answers up. In virtue of the

³ In a sample of 118 polar questions in the Switchboard Dialog Act Corpus ([Jurafsky, Shriberg & Biasca 1997](#)), 23 are superficially partial answers, and at least 5 of those employ the indirect strategy represented by (7). This seems like a healthy number in light of the fact that Switchboard participants were strangers talking on the phone, and thus had relatively little contextual knowledge to lean on.

way that these dialogues actually went, we will probably attribute some public commitments to the answerers, and this might have consequences for how we view their epistemic states, but the important work is done by the context.

However, not all indirect answers are as straightforward as these. Perhaps the best way to see that is to swing all the way from the natural cases above to more surprising dialogues like (8), which is a variation on an example that Matthew Stone discusses in his commentary.

- (8) A: Did John take a train from Paris to Istanbul?
 B: He likes spinach.

Let's join A in her desperate struggle to figure out what B means. We have learned that John likes spinach. By (3) and general principles of charitable interpretation, we will assume that B assumes that there is an inferential path from John's taste for spinach to an answer to A's question. Do we have access to some world knowledge that, when fed this knowledge, will deliver a resolving answer? Hard to say, hard to say. One thing is clear though: we are all of us (A especially) working overtime to try to find a way to relate B's reply to the immediate question concerning John's travel. As Stone says, "Our search for coherence demands an interpretive connection, rationality or no".

This is a source of speaker commitments to coherence relations, and thus it is fertile ground for the theoretical approach that A&L take in their work. We might never fully determine what B means to convey to us, and it might even be the case that B wishes to throw us off the scent, as it were, by forcing us to consider unimportant connections between spinach and rail travel. However, there is little doubt that, in this situation, B will end up committed to the following relationship (in the discourse situation of (8)):

- (9) *IQAP*(Did John take a train from Paris to Istanbul?, He likes spinach)

Two things about this are important:

- (10) a. People are attuned to content like (9). It is the sort of thing that we can debate, dispute, and query.
 b. Commitments of this form are robust across many kinds of discourse situation.

In their paper for the workshop, Franke, de Jager, and van Rooij discuss an example that brings together (10a, b). The example, from [Solan & Tiersma 2005](#), is drawn from a real courtroom situation (*Bronston v. United States*):

- (11) a. Prosecutor: Do you have any bank accounts in Swiss banks, Mr. Bronston?
 b. Bronston: No, sir.
 c. Prosecutor: Have you ever?
 d. Bronston: The company had an account there for about six months, in Zurich.

The literal content of Bronston's utterances is true: he did not have a Swiss bank account at the time of this conversation, and the company he mentions did in fact have one in Zurich for about six months. Bronston was judged to be dishonest nonetheless. Why? Because he too had once had a Swiss bank account. Though he never said anything that was literally inconsistent with this fact, he was perceived to have conveyed its negation. I think this dishonesty begins, at least, with an IQAP connection: in saying (11d) at this particular point in this particular discourse, he established an IQAP relationship with the question expressed in (11c):

- (12) *IQAP*((11c),(11d))

In the semantics of LoC, (12) has consequences for what Bronston believes about the relationship between his utterance and the question to which that utterance attaches. Now, he might try to wiggle out of this by saying that he intended (11d) to convey a 'Yes' answer. This will likely commit him to beliefs that he does not have — for example, an entailment between the company one works for having an account at bank *X* and having an account at *X* oneself.

Solan & Tiersma (2005) discuss situations like (11) at length, because they are important for establishing and arguing for subtle instances of perjury. Here is one of their constructed examples, again from an antagonistic environment:

- (13) John and Mary have recently started going together. Valentino is Mary's ex-boyfriend. One evening, John asks Mary, "Have you seen Valentino this week?" Mary answers, "Valentino's been sick with mononucleosis for the past two weeks." Valentino has in fact been sick with mononucleosis for the past two weeks, but it is also the case that Mary had a date with Valentino the night before.

Unless the date was very chaste, this deception is likely to catch up with Mary. When it does, Julio will likely think back on (13) and decide that Mary was dishonest. He is unlikely to say it in this way, but I think he would be within his rights to accuse Mary of falsely leading him to connect two utterances with *IQAP*.

What is important about the interactions involving Bronston and the prosecutor, and John and Mary, is that these are noncooperative scenarios. One speaker is trying to conceal information, and the other is trying to coax it out. Bronston and the prosecutor are even playing the sort of zero-sum game that A&L regard as epitomizing conflict. And yet, though this conflict was probably palpable through the proceedings, speakers still infer *IQAP* relations.

Thus, there is a tension between the claim that *IQAP* is rooted in cooperation and the empirically well-supported idea that it is a sort of automatic, form-based inference. I see two options. We could give up on the derivation from Cognitive *IQAP*, or we could give up on the notion that these dialogue games are not based on cooperativity at the level of communication.

These two approaches make different predictions. If *IQAP* is completely form-based, then it should remain powerful even in the most extreme cases of pure conflict. If it is rooted in cooperativity, then situations in which (2) become less powerful should also be situations in which speakers are reluctant to draw *IQAP* inferences. In such situations, they should insist on direct answers that leave little or no room for inference.

I think the weight of the evidence points to the idea that *IQAP* fades in especially antagonistic circumstances, even when the forms alone would seem to facilitate such connections. The recent example (14) is typical. Sheehan is Senator Norm Coleman's press agent. During the election run-up, Coleman was accused of accepting gifts from D.C. insiders and not reporting them in the manner demanded by Senate ethics guidelines. The reporters in this situation are trying to get Sheehan to respond to allegations that Coleman was given some fancy suits.⁴

(14) REPORTER: On a different subject is there a reason that the Senator won't say whether or not someone else bought some suits for him.

SHEEHAN: Rachel, the Senator has reported every gift he has ever received.

REPORTER: That wasn't my question, Cullen.

SHEEHAN: The Senator has reported every gift he has ever received. We are not going to respond to unnamed sources on a blog.

REPORTER: So Senator Coleman's friend has not bought these suits for him? Is that correct?

SHEEHAN: The Senator has reported every gift he has ever received.

⁴ The video is really something to see: <http://www.youtube.com/watch?v=VySnpLoaURL>. The video is titled 'A Simple Question, "Yes" or "No"??'

REPORTER: Why would say that? Why wouldn't you give us an answer yes or no on that?

SHEEHAN: The Senator has recorded every gift he has ever received.

REPORTER: We haven't asked whether he has recorded every gift he has ever received and I will take his word that he has recorded every gift he has received. Has he ever received a gift of suits?

SHEEHAN: The Senator has reported every gift he has ever received.

[Sheehan says 'The Senator has reported every gift he has ever received' five more times in two minutes.]

These reporters systematically refuse to regard Sheehan's "The Senator has reported every gift he has ever received" as responsive to their questions, though this is not inherently different from what happens in (4) and (7).

A&L are well aware that the relevant notions of cooperativity are variable in strength and focus. One sees this in their qualifications to the *Grice Cooperative (GC)* element (p. 25). This snippet from a recent email from Alex makes this even clearer:

cooperativity is a continuum rather than a binary class, and in many contexts where non-cooperative things are going on people **appear** to be behaving cooperatively (and doing so in order to achieve their 'non-cooperative' goals).

I think this is exactly right. I also think it is exactly right to say, as LoC does, that Cooperativity lies at the heart of IQAP. My conclusion, then, is that there is a larger role for cooperation than the rhetoric in 'Commitments' would lead one to expect.

4 Beyond IQAP: The nature of enrichment

In the discussion of indirect answers in the previous section, there were really just two ingredients: (i) the discourse coherence relations (I talked only about *IQAP*, but others are relevant), and (ii) inferences from world knowledge, including knowledge of the context.

Much of the literature on Gricean pragmatics could lead one to believe that this picture is incomplete — that we still need to sort out the role that the maxims play in these inferences. I am not sure, though, that there is much more to be said. I do not know how to make a convincing case for this sweeping statement, but I can report on some small experiments that might at least convey why I am willing to stake out this position.

The experiments (simulations, really) are conducted using the software suite *Alchemy* (Kok, Sumner, Richardson et al. 2006), which was developed by Pedro Domingos and his colleagues at the University of Washington.⁵ *Alchemy* is a powerful tool for doing a variety of tasks in machine learning. I think it wasn't designed to do pragmatic analysis, but it is suited to it nonetheless.

In general, the set-up works like this: one writes down a series of constraints in first-order logic. Hard constraints are finished with a period. Soft (defeasible) constraints are preceded by a real-valued weight and lack a period. Here is an example adapted from Richardson & Domingos (2006):

```
(15) // Friends of friends are friends.
      0.7 Friend(x, y) ^ Friend(y, z) => Friend(x, z)
      // Smokers have cancer.
      1.5 Smoke(x) => Cancer(x)
      // Friends have the same smoking habits.
      1.0 Friend(x, y) => (Smoke(x) <=> Smoke(y))
      // Rigid facts about the world.
      Friend(Ann, Bob)      Friend(Bob, Cal)
      Smoke(Ann)            Smoke(Bob)
```

Ideally, we would learn these weights from data; *Alchemy* has functionality for doing just that. Since I am not entirely sure what the training data would need to be like in the scenarios below, I just stipulate the weights throughout this commentary.

The situation described in (15) does not contain any direct factual statements about who has cancer. Nonetheless, it has the right stuff for doing inference about this predicate. If we ask *Alchemy* `?Cancer(Ann)`, it return a relatively high probability: around 0.8. As we increase the weight on `Smoke(x) => Cancer(x)`, this confidence measure increases; in the limit (if this is changed to a hard constraint), `?Cancer(Ann)` is answered positively with a probability of essentially 1.

Alchemy has efficient methods for searching large spaces of possible worlds using Monte Carlo methods, so one can make these context files very rich. We can test intricate situations while at the same time keeping track of all the assumptions that go into the inferential behavior. The limitation to first-order logic can be a drag, but the domains are all finite, so it is easy to simulate second-order expressivity where it is needed (for belief statements, for specifying what the question under discussion is, and so forth).

In the previous section, I asserted that it is straightforward to ensure that

⁵ The software and documentation are at <http://alchemy.cs.washington.edu/>.

Did John pass? is completely answered by *He got 60%*. Here is an Alchemy implementation, using only hard constraints, that delivers exactly that result:

```
(16)  A contextual fragment for (4)

      // Generalization.
      Score(x,Sixty) => Pass(x).

      // Fact about the world.
      Score(John, Sixty)
```

In this context, the chances of John having passed are estimated to be (essentially) 100%. However, if you are not totally sure of $\text{Score}(x, \text{Sixty}) \Rightarrow \text{Pass}(x)$, then you might lower the weight on this generalization. This will inject uncertainty into the extent to which this reply resolves the issue.

With this background in place, let's turn to a case that naturally creates more uncertainty about the communicative value of the answer:

```
(17)  A: Is Sue at work today?
      B: She is sick with the flu.
```

Just what is B saying with this reply? In cooperative situations, A might perceive B as having given a “No” answer. Then again, suppose that B intends this to be a complaint about how Sue is going to infect the office. Or suppose that B intends to mislead, getting Sue into trouble unfairly (as we saw with Bronston and Mary above).

Different background assumptions deliver different values in these cases. The generalizations in (18) are intended to capture some general ideas about how we regard working, being sick, and the connections between those properties. None could safely be considered a hard constraint, but all have a glimmer of truth, so we can throw them all in.

```
(18)  a. Sick(x) => !AtWork(x)
      b. HardWorking(x) => AtWork(x)
      c. (Malicious ^ Sick(x)) => AtWork(x)
      d. (AtWork(x) ^ Sick(x)) =>
          (Malicious(x) v Thoughtless(x))
```

In Alchemy, $!\varphi$ means ‘not φ ’. Thus, if (18a) is a hard constraint (or, very heavily weighted), and we assume that B is truthful, then we achieve a high degree of confidence that Sue has stayed home:

```
(19) 20.0 Sick(x) => !AtWork(x).
      2.0 HardWorking(x) => AtWork(x)
      1.0 (Malicious ^ Sick(x)) => AtWork(x)
      1.0 (AtWork(x) ^ Sick(x)) =>
            (Malicious(x) v Thoughtless(x))

// Facts.
Sick(Sue)    HardWorking(Sue)    !Malicious(Sue)
Inference: AtWork(Sue): 0.00005
```

However, if we regard all of these pieces of information as more or less equally weighted, then Sue's being hard-working dramatically raises the probability that she is at work:

```
(20) 1.0 Sick(x) => !AtWork(x)
      1.0 HardWorking(x) => AtWork(x)
      1.0 (Malicious ^ Sick(x)) => AtWork(x)
      0.5 (AtWork(x) ^ Sick(x)) =>
            (Malicious(x) v Thoughtless(x))

// Facts.
Sick(Sue)    HardWorking(Sue)    !Malicious(Sue)
Inference: AtWork(Sue): 0.439906
```

As we add contextual information, the inferences grow richer. I have not, for example, included factors pertaining to relative trust levels, possible ulterior motives, and so forth, nor have I implemented the coherence relations themselves. I see no obstacles to including such constraints, though. My sense is that implicature behavior, of the sort often attributed to the conversational maxims, will emerge from such simulations.

5 Other issues we could place under discussion

The central argument of the above commentary is that cooperation is important even in antagonistic discourse relations. My primary evidence is that the discourse coherence relation *IQAP*, which is rooted in cooperativity, remains powerful even when nonlinguistic preferences are misaligned. In addition, section 4 argues briefly that coherence relations and general inferences from contextual information can account for a lot of what falls under the heading of conversational implicature.

Much more could be said, though, about A&L's provocative, insightful body of work. I don't have the time to say much more right now, but I do want to

close by mentioning two issues that are near and dear to me and that A&L's work can inform in many ways.

Discourse particles These seem to provide lexical evidence for various discourse relations. They can signal incongruence, surprise, explanation, and a host of other relationships. Though they are generally optional, discourse conditions can make them essentially obligatory. There is great potential here for further blurring the distinction between conventionalized meaning and pragmatic enrichment.

Expressives Expressive language — swears, honorifics, epithets, exclamatives — takes its toll on the people who hear it, both psychologically and physiologically (Jay 2000; Jay, Caldwell-Harris & King 2008). However, it can also provide an important release for the speaker, and it can be used strategically, to make content more memorable, powerful, or relevant (Constant, Davis, Potts et al. 2008). This seems like ripe territory for signaling games of the sort explored by van Rooy 2003. I can show that expressive signaling systems are stable despite potentially very high costs to the speaker. It would be fruitful to consider such systems in the rich, layered logical setting of A&L's work.

References

- Asher, Nicholas & Alex Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.
- Asher, Nicholas & Alex Lascarides. 2008. Making the right commitments in dialogue. Paper presented at the University of Michigan Linguistics and Philosophy Workshop, November 21-23.
- Büring, Daniel. 1999. Topic. In Peter Bosch & Rob van der Sandt (eds.) *Focus — Linguistic, Cognitive, and Computational Perspectives*, 142–165. Cambridge: Cambridge University Press.
- Constant, Noah, Christopher Davis, Christopher Potts & Florian Schwarz. 2008. The pragmatics of expressive content: Evidence from large corpora. To appear in *Sprache und Datenverarbeitung*.
- Groenendijk, Jeroen & Martin Stokhof. 1984. *Studies in the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, University of Amsterdam.
- Jay, Timothy. 2000. *Why We Curse: A Neuro-Psycho-Social Theory of Speech*. Philadelphia/Amsterdam: John Benjamins.
- Jay, Timothy, Catherine Caldwell-Harris & Krista King. 2008. Recalling taboo and nontaboo words. *American Journal of Psychology* 121(1): 83–103.

- Jurafsky, Daniel, Elizabeth Shriberg & Debra Biasca. 1997. Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual, draft 13. Tech. Rep. 97-02, University of Colorado, Boulder Institute of Cognitive Science, Boulder, CO.
- Kok, Stanley, Marc Sumner, Matthew Richardson, Parag Singla, Hoifung Poon & Pedro Domingos. 2006. The Alchemy system for statistical relational AI. Tech. rep., Department of Computer Science and Engineering, University of Washington, Seattle, WA.
- Richardson, Matthew & Pedro Domingos. 2006. Markov logic networks. *Machine Learning* 62(1–2): 107–136.
- Roberts, Craige. 1996. Information structure: Towards an integrated formal theory of pragmatics. In Jae Hak Yoon & Andreas Kathol (eds.) *OSU Working Papers in Linguistics*, vol. 49: Papers in Semantics, 91–136. Columbus, OH: The Ohio State University Department of Linguistics. Revised 1998.
- van Rooy, Robert. 2003. Being polite is a handicap: Towards a game theoretical analysis of polite linguistic behavior. In Moshe Tennenholtz (ed.) *Proceedings of TARK 9*, 45–58. TARK. URL http://www.tark.org/proceedings/tark_jun20_03/proceedings.html.
- van Rooy, Robert. 2004. Signalling games select Horn strategies. *Linguistics and Philosophy* 27(4): 493–527.
- Solan, Lawrence M. & Peter M. Tiersma. 2005. *Speaking of Crime: The Language of Criminal Justice*. Chicago, IL: University of Chicago Press.