What Game Theory can do for NLG: the case of vague language

Kees van Deemter

University of Aberdeen

k.vdeemter@abdn.ac.uk

Abstract

This informal position paper brings together some recent developments in formal semantics and pragmatics to argue that the discipline of *Game Theory* is well placed to become the theoretical backbone of Natural Language Generation. To demonstrate some of the strengths and weaknesses of the Game-Theoretical approach, we focus on the utility of *vague* expressions. More specifically, we ask what light Game Theory can shed on the question when an NLG system should generate vague language.

1 NLG as a choice problem

Natural Language Generation (NLG) is the area of computational linguistics that is concerned with the mapping from non-linguistic to linguistic expressions (e.g. Reiter and Dale 2000). This formulation might be taken to suggest that NLG is best viewed as a kind of *translation* problem, where the challenge is to find a way to convert a formal expression into (for example) an English one. In its early years, this may have been a fruitful way to think about NLG but, these days, a better perspective is of NLG as a *choice* problem. For after the advances of recent years, the problem is no longer such much “How on Earth can this information be expressed in English?”, but rather “From all the possible ways to express this information in English, which one is the most effective choice?”

Let us try to say this a bit more precisely. It is usually fair to assume that the formal expressions from which NLG takes its departure are themselves clear and unambiguous. Let us call the inputs to the generator *Meanings*. Now suppose we have a grammar that tells us how each given Meaning can be expressed in a language such as English. The task for NLG now is to choose, for

each of these Meanings, which of all the different linguistic Forms that can express it (according to the grammar) is the *best* expression of this particular Meaning. Ultimately, this choice is likely to depend on a number of other parameters, such as the identity of the hearer, and the words that have earlier been used. In the present paper, these “contextual” issues will largely be ignored, allowing us to simplify by thinking in terms of a mapping from Meanings to Forms.

The perspective that views NLG as a choice problem is far from new (see e.g. McDonald 1987, where it takes a central position); in fact, it forms the methodological spine of Systemic-Functional Grammar, with its AND/OR graphs (Bateman 1997). Given this perspective, the question comes up what factors determine the choice between different linguistic Forms. This question is difficult to answer in detail, but at the most abstract level, the answer is likely to have something to do with the “utility” of the different Forms that can be generated, and perhaps such additional factors as the *cost* to the speaker of generating them, and the *cost* to the hearer of processing (e.g., parsing and interpreting) them. To utter a sentence is to perform an action, and the choice between different actions is naturally thought of as governed by utility, understood in the broadest possible sense.

2 Game Theory

The analysis of NLG as driven by the *utility* of utterances feels natural to people familiar with practical applications of NLG, where texts are generated for a real-life setting. More generally, this type of analysis suits anyone who is interested in the effects of an utterance on an audience (e.g., Mellish and Van der Sluis 2009). To see how NLG systems could be amenable to a decision-theoretical analysis, in which the expected pay-offs associated with different texts are compared, consider an NLG system that informs roadgrit-

ters' decisions about the condition of the roads in Scotland, to help them decide which ones are icy enough to require treatment (e.g. Turner et al. 2008).

Computerised weather forecasts can tell road gritters which roads are likely to be icy, and hence dangerous. There can be thousands of dangerous roads on a given night, and it is often impossible to say in a few words *exactly* which road are dangerous (Turner et al. 2008). One summary produced by the generator might approximate the data by saying 'Roads in the Highlands are icy' while another might say 'Roads above 500 metres are icy' (assume this covers a larger set of roads). It matters which of these summaries is generated, because each summary will lead to a different set of roads being treated with salt (i.e., gritted). The first summary may have 10 false positives (i.e., roads gritted unnecessarily) and 10 false negatives (i.e., dangerous roads not gritted); the second summary might have 100 false positives and only 2 false negatives. In a situation of this kind, which involves a tradeoff between safety on the one hand, and money and environmental damage (from salt) on the other, decision theory would be a natural framework in which to compare the utility of the two summaries. If a false positive has a negative utility of -0.1 and a false negative one of -0.5 , for example, then the first summary wins the day. (Needless to say, the choice of these constants is crucial, and tricky to justify.)

More specifically, many NLG systems invite a *game-theoretical* analysis – or an Optimality-Theoretic analysis, which can come down to the same thing (Dekker and Van Rooij 2000; Van Deemter 2004 for an application to NLG). Suppose I want to state that all old people are entitled to certain benefits (cf. Khan et al. 2009):

- a. Old men and old women are entitled to benefits.
- b. Old men and women are entitled to benefits.

Which of these two linguistic Forms should I choose? This depends on the strategy of the hearer. If the hearer interprets (b) as concerning *all* women (rather only the old ones) then my utterance will have misfired to an extent. The success (for speaker and/or hearer!) of the speaker's generation strategy, in other words, depends on the hearer's interpretation strategy.¹

¹For a game-theoretical perspective on the generation of

This interaction means that decision theory is not the best tool for analysing the situation, for wherever different agents' strategies interact, *Decision Theory* gives way to *Game Theory*. Game Theory was conceived in the nineteen forties (Von Neumann and Morgenstern 1944) and has since come to be used extensively by economists, sociologists, biologists, and others. Far from being limited to games in a limited sense of the word, Game Theory is the mathematical study of rational social interaction and, as such, it is reasonable to expect it to be able to shed light on language use as well. Perhaps more than anything, it promises to have the potential to explain *why* communication works the way it does. For if we could show that people's linguistic behaviour conforms with what it would be rational for them to do, then this would have substantial explanatory value.

Work by David Lewis and other on communication and *coordination games* helped to make Game Theory relevant for situations where the players are not in conflict with each other (Lewis 1969). A classic example is where two generals are both intent on attacking an enemy, but while each general individually is weaker than the enemy, they can beat him if they attack at the same time. Communication ("I am going to attack now!") can help the generals to cooperate and win the battle. Essentially the same things happens when you try to meet a friend: neither of you may care where and when to meet, as long as the two of you end up in the same place at the same time; communication, of course, can help you achieve this goal.

Applications of Game Theory to language now come in many flavours (see e.g. Klabunde 2009, this conference). In this paper I want to engage in a small case study: the expression of *quantitative information* in English. More specifically, I will focus on the fact that quantitative information is often only communicated *vaguely*. When a thermometer, for example, measures your body temperature as 39.82 Celcius, your doctor might express this by saying that your temperature is '39.8 degrees', but he might also round this off even further, saying that it is 'approximately 40 degrees'. Even more vaguely, she might tell you that you have 'a high fever'. Which of these linguistic Forms is preferable, and why?

Questions of this kind have led to a lively dis-

referring expressions, where success depends on alignment between hearer and speaker strategies, see Kibble 2003.

cussion among linguists, philosophers, and theoretical economists (Lipman 2000, 2006; De Jaegher 2003; Van Rooij 2003), focussing on the question under what circumstances vagueness can lead to a higher utility than crispness. The question is important for understanding human communication, because vagueness plays such a central role in it. Vague adjectives, for example, are prevalent among the words first learned by a typical infant (Peccei 1994) and many of their subtleties are understood by children of only 24 months old (Ebeling and Gelman 1994). In my opinion, the understanding of vagueness is equally important for the NLG community, and particularly for those of us who work on “data to speech” (Theune 2001) or “data to text” (Reiter 2008) systems, where the expression of quantitative data plays such a crucial role. For this reason, I have chosen it as the topic of an informal case study on the relevance of game theory for NLG.

3 Vagueness in situations of conflict

First, let us focus on a type of situations where it is relatively easy to understand what the differential benefits of vagueness can be. We start by examining a game-theoretic study of a different phenomenon: ambiguity.

3.1 The utility of ambiguity: Aragonès and Neeman

Like others who have discussed these issues, we take *vagueness* to arise if an expression allows borderline cases. The word ‘tall’, for example, is vague because in a typical context there can be people who are difficult to categorise as either tall or not tall: they are somewhat in between, one is tempted to say. *Ambiguity* is something else. It arises when an expression can be meant in a limited number of different ways. The word ‘letter’, for example, is ambiguous because it can refer to one individual character or to the sort of meaningful arrangement of characters that people once used to communicate long distance. In 1994, two game theorists asked whether a Game Theoretical explanation might be given for strategic use of ambiguity (i.e., where ambiguity is used on purpose), and they came up with the following answer (Aragonès and Neeman 1994).

Suppose two unscrupulous politicians position themselves for an election. Not burdened with any convictions, they are free to choose between

three different ideologies (left, right, center), depending on what gives them the highest utility; additionally, they can choose between two *commitment*² levels, c_{high} and c_{low} , both representable as real numbers with $c_{high} > c_{low}$. Unfortunately, Aragonès and Neeman do not say what a commitment level is, but one might think of a more and a less extreme version of their chosen ideology.

What combination of an ideology and an commitment level should each politician choose? This depends on the electorate, of course. Suppose there are three blocs of voters: V(left), V(right) and V(center). A leftist voter prefers a leftist politician, and preferably one with a high commitment level. Confronted with a choice between two rightwing politicians, our leftist voter will prefer one with a low commitment. A rightwing voter behaves as the mirror-image of the leftist voter, while the neutral voter is neutral between the two ideologies but, weary of ideology, she prefers low commitment over high commitment. Commitment, in other words, is only relevant for a choice between politicians of the same ideology.

If this is the whole story then politicians will choose an ideology and commitment level based on their estimates of the numbers of voters in each bloc, trying to maximise their expected payoff, formulated solely in terms of the likelihood of winning the election. The task for Game Theory is to work out what *combination* of strategies might give both politicians the highest possible payoff, for example in the sense that a policy change by just one of the two politicians can never improve his expected payoff.

But Aragonès and Neeman’s model allows politicians to look beyond the election, towards their anticipated time in government. Surely, a low commitment is easier to fulfil than a high commitment, particularly in view of unforeseen contingencies, so it is nicer to be elected on a low-commitment platform that does not tie one’s hands too much. To model this, Aragonès and Neeman formulate utility in a way that multiplies the probability of a politician’s winning the elections with a constant that is negatively correlated to his commitment. Let $U_i(I_1, c_1; I_2, c_2)$ be the utility for politician i given that politician 1 chooses ideology I_1 with commitment level c_1 ,

²Aragonès and Neeman call these *ambiguity* levels, but since the relation with ambiguity is debatable we opt for a more neutral term. Low commitment equals high ambiguity and conversely.

while politician 2 chooses I_2 with level c_2 . Furthermore, $P_i(I_1, c_1; I_2, c_2)$ represents the probability of i winning the elections given this same constellation of choices. Let $k \geq c_{high}$.

$$\text{Utility formula: } U_i(I_1, c_1; I_2, c_2) = P_i(I_1, c_1; I_2, c_2)(k - c_i)$$

Under these assumptions one can show that a low commitment level (i.e., c_{low}) can sometimes give a politician a slightly *lower* probability of winning the elections (because his core voters will be less inclined to vote for him), yet a *higher* overall utility (because his time in office will be easier). For details see Aragonès and Neeman (1994).

3.2 The utility of vagueness

It is often thought that Aragonès and Neeman's model demonstrates how ambiguity can be used strategically, but that it fails to shed light on *vagueness* (e.g. De Jaegher 2003). I do not see, however, how this view stands up to linguistic scrutiny. To see why, let me construct what strikes me as a possible example.

Suppose the ideology in question – a leftist, or perhaps a populist one – is to take away money from the 10% of richest people and give it to the 10% poorest. Commitment level, in this case, could be a way of making explicit *what percentage* of the top 10% to give away. One position might assert that this has to be, say, 50% of their income, while another position might put this figure at 5%. But if we identified high commitment with the 50% position and low commitment with the 5% position then none of the two commitment levels would be ambiguous. To make one of them ambiguous, we would need something like the following:

- The ambiguous politicians' game:
- I : take money from the 10% of richest people and divide it equally over the 10% poorest.
 - c_{50} : do I with 50% of the money of each of the richest people.
 - $c_{ambiguous}$: do I with either 5% or 50% of the money of each of the richest people.

But this must be a simplification, for we are dealing with a continuum: there is nothing to exclude percentages in between 25% and 5%, for example. It seems, therefore, perfectly possible to construct a version of Aragonès and Neeman's game

– an even more plausible version, I believe – that hinges on vagueness. For example:

- The vague politicians' game:
- I and c_{50} : (as above).
 - c_{vague} : do I with a *large* portion of the money of each of the richest people

Clearly, c_{vague} involves vagueness, because 'a large portion' admits borderline cases. In all important respects the vague politicians' game is isomorphic to the ambiguous politicians' game: fierce advocates of redistribution would favour c_{50} over c_{vague} , for example, because the latter leaves them uncertain over the amount of redistribution. It is also plausible that politicians would prefer to avoid a commitment as clear as c_{50} , because future contingencies might make it difficult for them to honour this promise. In fact, one could extend the game with a second election, in which the electorate could give their verdict on a politician's time in office, and to adapt the utility formula with a third term which represents the probability of winning that second election. Surely, the breaking of promises doesn't do much for a politician's changes of being re-elected, and a precise promise is easier to break than a vague one.

With help from Aragonès and Neeman, we have found a situation in which vagueness has a higher utility than precision.³ It should be noted, however, that this model (and that of De Jaegher 2003 likewise) hinges on the fact that the interests of the speaker and the hearer differ: what's good for the politician may be bad for his voters. NLG systems can be faced with similar asymmetries, for example when an artificial doctor decides to keep its predictions vague to avoid being contradicted by the facts; a doctor who says "These symptoms will disappear fairly soon" is less likely to get complaints than one who says "These symptoms will have disappeared by midnight next Sunday". Something similar holds for a roadgritting system (like the one in Turner et al. 2008), which might easily face lawsuits if it gets things too evidently wrong. Advertisements also come to mind, be-

³Another game with this property was described in De Jaegher (2003), involving a more complex version of the game of the *two generals* (section 2). De Jaegher's game lets one general tell the other about the *preparedness* of the enemy. The utility of vagueness hinges on a subtle asymmetry between the generals, only one of whom will suffer if the enemy turns out to be *prepared*. Intriguing though it is, I find it difficult to see how De Jaegher's game is relevant to everyday communication or NLG.

cause the interests of the advertiser may not coincide with those of the customer. – Examples where vagueness can save money or face are plentiful, yet one wonders whether vagueness can also be advantageous in situations where it is one’s honest aim to inform an audience as well as one can.

4 Vagueness when there is no conflict

So, let us investigate the advantages of vagueness in situations that are typical for today’s NLG systems, where the system tries, unselfishly, to assist a user to the best of its ability.

4.1 Lipman’s questions

The question why vagueness is used strategically in situations where the interests of speakers and hearers are essentially aligned was asked perhaps most forcefully by the economist Barton Lipman. First he did this in a brief response to an essay by the famous game theorist Ariel Rubinstein (Lipman 2000), and later in a growing but still unfinished discussion paper (Lipman 2006). Lipman uses what we shall call an *airport scenario*, where player 1 asks player 2 to go to the airport to pick up an acquaintance of player 1. In its simplest form, the scenario lets player 1 know the referent’s height with full precision (assuming that such a thing is possible), while player 2 carries a perfect measuring device. There are two other people at the airport, and it is assumed that heights are distributed uniformly between a maximum denoted by 1 and a minimum denoted by 0. The payoff for both players – please note the symmetry! – is 1 if player 2 successfully picks the referent, while it is 0 if she fails (i.e., the first person she addresses turns out to be someone else).

Lipman observes that, under these assumptions, vagueness would be bad: why would player 2 say ‘He is tall’, for example, if he can say ‘He is 183.721cm’? By stating his acquaintance’s exact height, player 1 will allow player 2 to identify this person with almost complete certainty, given that the chance of two people having the exact same height is almost nil. Lipman also wonders what would happen if only one predicate was available to player 1. He proves that, under these assumptions, optimal communication arises if a word is used in accordance with the following rule:

Say ‘the tall person’ if $\text{height}(\text{person}) > 1/2$, else say ‘the short person’.

Lipman observes that this concept of ‘tall’ does not involve vagueness, because the rule does not allow any borderline cases: everyone is either tall or short. In other words, no *rationale* for vagueness has yet been found.

Note that Lipman is not questioning that vague utterances can be useful, which they evidently can be (see e.g. Parikh 1994 for a convincing demonstration using a game-theoretic approach). He is asking whether vague expressions can be *more* useful than any non-vague expression.

4.2 Answering Lipman

First, let us consider a possible modification of Lipman’s scenario. In this *modified airport scenario* the speaker knows the heights of all three people at the airport. Suddenly it becomes easier to understand why vagueness can be useful. For suppose your acquaintance happens to be the tallest person of the three. You can then identify him as ‘the tall guy’. Arguably, this is safer than citing the person’s height in centimeters, because ‘the tall guy’ (meaning, in this case, the same as *tallest* guy) does not require the players to make any measurements: *comparisons* between heights can often be made in an instant, and with more confidence than absolute measurements. I dealt with cases of this type in my paper on vague descriptions (van Deemter 2006), where a generator takes numerical height measurements to produce noun phrases that involve gradable adjectives: ‘the tall guy’, ‘the fastest one of the three heavy tortoises’. In cases like this, one can argue that vagueness is only *local*, in the same way that ambiguity can be merely local, for example when the sentence as a whole allows one to disambiguate an ambiguous word in it (e.g. when a pronoun gets resolved or a lexical item disambiguated). In the modified airport scenario, the noun phrase as a whole (e.g., ‘the tall guy’) allows no borderline cases, so there is no *global* vagueness here.

Local vagueness is wide-spread and can make use of different “precisification” mechanisms. When I say of a gymnastic exercise, for example, that it is ‘good for young and old’, for example, then there is nothing vague about my description of the people involved: I am using vague words to say that this exercise is good for *everyone*, regardless of age. Although local vagueness constitutes some kind of answer to Lipman, most linguists assume that *globally* vague utterances exist

as well (even when the interests of the speaker and the hearer are aligned). Let us assume they are right and continue to look for a *rationale*.

Secondly, it has been suggested that strategic vagueness can arise from a desire to reduce the “cost” of the utterances involved (e.g. Van Rooij 2003, Jäger 2008). One might amplify this idea by arguing that vague words are part of a highly efficient mechanism that makes their meaning dependent on the context in which they are used. The size constraints on ‘a small elephant’, for example, are very different from those on ‘a small mouse’; this suggests that vague words may not only be efficient to use but also efficient to learn (Van Deemter, in preparation). All this seems true enough but, as an answer to the question “Why vagueness?” it does not stand up to Lipman-style scrutiny. Let me explain why not.

Consider the earlier-mentioned doctor, who measures your body temperature as 39.82 Celcius. By stating that you have ‘a high fever’ (instead of ‘thirty eight point eighty two degrees’) the doctor is pruning away details that are of questionable relevance in the situation at hand. But this does not force him to use language that is *vague*: language that allows borderline cases, in other words. He could have achieved a similar economy by rounding, saying that your temperature is ‘(about) forty degrees’; in this way, he would have reduced information without being vague. The benefits of information reduction can be modelled in a game where communication informs action: if ‘38.82 Celcius’ and ‘39 Celcius’ are associated with the same medical action (e.g., to take an aspirin) then the fact that ‘39 Celcius’ is “cheaper” to produce and to process will tend to give this expression a better utility than ‘38.82 Celcius’. But information reduction does not imply vagueness, so we are back at square one: Why vagueness?

It might be thought that things change when uncertainty is taken into account: a measurement of 39.82 Celcius is not as exact as it sounds, for example, because errors are likely. The result of the measurement is perhaps best conveyed by a normal distribution of which 39.82 is the mean value, and such a complex curve is difficult to put in just a few words. Still, the argument of the previous paragraph continues to apply, because the curve can be summarised without vagueness: the figure of 38.82 Celcius is one such summary.

A *third* suggestion (e.g. Veltman 2002) is that

vague expressions such as ‘high fever’ do more than just *reduce* the information obtained from a measurement. The expression ‘high fever’ also *adds* bias or evaluation to the raw data, namely the information that the temperature in question is worrisome. You do not need domain knowledge to understand the medical implications: hearing that something medical is ‘high’ tells you that you should be worried.

Once again, this sounds like an excellent reason for using vagueness, particularly in situations where an understanding of the metric in question cannot be taken for granted (such as oxygen saturation, the metrics for which mean little to most of us). Still, it is not evident that this justifies the use of vagueness. If bias needs to be expressed, then why not simply add it? Why not state the exact temperature (or an approximation of it) *and* say that this reading should be considered worrisome? One might respond that this would have been time and space consuming, but if that is a problem then why have no conventions arisen for expressing quantities in two ways, a worrisome and a non-worrisome one? Why should bias necessarily be coupled with vagueness only, given that it is as easy to think of a crisp expression that contains bias as it is to think of a vague expression that does not (e.g., in the case of an adjective like ‘tall’)? A good example of *crispness* + *bias* is the word ‘obese’, in the sense of having a Body Mass Index of over 30. For the reason why obesity was defined in this way is precisely that this degree of overweight is considered medically worrisome.

5 Discussion: vagueness and game theory

5.1 Vagueness is harder to justify than you think

Let us first summarise our findings about vagueness, some of which will be discussed more fully in Van Deemter (in preparation).

It is often easy to see why vague *words* come in handy; the modified airport scenario demonstrates how vague words can create an information loss that is only local: by making a vague word part of a referring expression, a crisp borderline is enforced on a vague concept, resulting in a beautifully efficient description (e.g., ‘the tall guy’) that is arguably clearer than any expression that relies on absolute values. This means that the utterance *as a whole* is not vague at all: it is only locally vague. Whether we speak of vagueness in such

situations seems a matter of taste.

It is also clear why vagueness can have differential benefits in communication between agents whose interests differ more than just minimally (cf., Horton and Keysar 1996 for experimental evidence of speaker's laziness in situations where their interests are approximately aligned), such as a politician and his potential voters, or like a professional who does not wish to be sued by his clients. In situations of this kind it can be beneficial for a speaker or an NLG system to obfuscate, exploiting the borderline cases inherent in vague expressions.

Beyond this, it is surprisingly difficult to see how vagueness can be advantageous for NLG. This is partly because there appear to exist some linguistic issues that NLG researchers are able to disregard. It seems plausible, for example, that vagueness is unavoidable in situations where no commonly understood metrics are available, for instance when we judge how beautiful a sunset is, how wise a person, or how dismal the weather. As long as NLG systems use tangible input data (about millimetres rainfall, for example, or body temperature), these reasons for vagueness seem irrelevant. Similarly, there is much that is unknown about the working of perception even in simple domains. What is it that allows me to talk about the height of someone I see, for example? The input to my personal "generator" (as opposed to the input to a typical NLG system) might not be equivalent to a tidy number. (Could it be some inherently vague percept, perhaps?) These difficult questions (see also Lipman 2006) must remain unanswered here.

5.2 The utility of utility

Confronted with the claim that Game Theory should be the theoretical backbone to NLG, some people might respond that no new backbone is needed, because the theory of formal languages, conjoined with a properly expressive variant of Symbolic Logic, provides sufficient backbone already. I believe this objection to be misguided. Admittedly, the disciplines in question are well suited for saying which Forms can express which Meanings. But it is far less clear that these disciplines have anything to say about the key problem of NLG: how to choose the most effective way to express a given Meaning in (for example) English. This is a vacancy that Game Theory would be well placed to fulfill, in my opinion. The present paper

has illustrated this claim by discussing the question when and why a generator should choose a vague expression. The fact that this discussion has yet to produce a clear conclusion is, in my opinion, not due to any shortcomings of Game Theory, but to the intrinsic difficulty of the problem.

There is, of course, a *caveat*. The use of game theory in empirical sciences has, with proper modesty, been described as "modelling by example" (e.g. Rasmussen 2001): a mathematical game shows us an example of how things *might* be, not necessarily how things are. The situation is familiar to linguists, of course, and from applications of mathematics more generally. By inspecting a formal grammar, for example, one does not learn much about language, unless there exists evidence that the linguistic Forms and Meanings pair up as specified by the grammar. In similar fashion, one learns little from a Game Theoretical model unless one has reason to accept the assumptions that were built into it: the choices that it assumes available to the players, and the payoffs related to each outcome of the game, for example. This means that Game Theory can come to the aid of linguistic pragmatics and NLG, but that only empirical research can tell us what games people *actually* play when they communicate.

Acknowledgments

Thanks are due to my colleagues Ehud Reiter, Albert Gatt and Hans van Ditmarsch, for useful discussions on the theme of this paper. Funding from the EPSRC under the Platform Grant "Affecting People with Natural Language" (EP/E011764/1) is gratefully acknowledged.

References

- Aragonès and Neeman 2000. Enriqueta Aragonès and Zvika Neeman. Strategic ambiguity in electoral competition. *Journal of Theoretical Politics* **12**, pp.183-204.
- Bateman 1997. John Bateman. Sentence generation and systemic grammar: an introduction. Iwanami Lecture Series: Language Sciences. Iwanami Shoten Publishers, Tokyo.
- de Jaegher 2003. Kris de Jaegher. A game-theoretical rationale for vagueness. *Linguistics and Philosophy* **26**: pp.637-659.
- Dekker and Van Rooij 2000. Bi-directional Optimality Theory: an application of Game Theory.

Journal of Semantics **17**: 217-242.

Ebeling and Gelman 1994. K.S.Ebeling and S.A.Gelman. Children's use of context in interpreting "big" and "little". *Child Development* **65** (4): 1178-1192.

Horton and Keysar 1996. William S. Horton and Boaz Keysar. When do speakers take into account common ground? *Cognition* **59**, pp.91-117.

Jäger 2008. Gerhard Jäger. Applications of Game Theory in Linguistics. *Language and Linguistics Compass* **2/3**.

Khan et al 2009. Imtiaz Khan, Kees van Deemter, Graeme Ritchie, Albert Gatt, and Alexandra A.Cleland. A hearer oriented evaluation of referring expression generation. Proc. of 12th European Workshop on Natural Language Generation (ENLG-2009).

Kibble 2003. Rodger Kibble. Both sides now: predictive reference resolution in generation and resolution. Proc. of Fifth International Workshop on Computational Semantics (IWCS-2003). Tilburg, The Netherlands.

Klabunde 2009. Ralph Klabunde. Towards a game-theoretic approach to content determination. Proc. of 12th European Workshop on Natural Language Generation (ENLG-2009).

Lewis 1969. David Lewis. *Convention – A Philosophical Study*. Harvard University Press.

Lipman 2000. Barton L.Lipman. Economics and Language. "Comments" section, Rubinstein (2000).

Lipman 2006. Barton L.Lipman. Why is language vague? Working paper, December 2006, Department of Economics, Boston University.

McDonald 1987. Natural Language Generation. In S.Shapiro *Encyclopaedia of Artificial Intelligence*, Volume 1. John Wiley, New York.

Mellish and Van der Sluis 2009. Chris Mellish and Ielka van der Sluis. Towards empirical evaluation of affective tactical NLG. Proc. of 12th European Workshop on Natural Language Generation (ENLG-2009)

Von Neumann and Morgenstern 1944. John von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Wiley & Sons, Princeton, New Jersey.

Parikh 1994. Rohit Parikh. Vagueness and utility:

the semantics of common nouns. *Linguistics and Philosophy* **17**: 521-535.

Peccei 1994. Jean Stilwell Peccei. *Child Language*. Routledge.

Rasmussen 2001. Eric Rasmussen. *Games & Information: an introduction to game theory*. Third Edition. Blackwell Publishing.

Reiter and Dale 2000. Ehud Reiter and Robert Dale. *Building natural language generation systems*. Cambridge University Press. Cambridge.

Reiter 2007. Ehud Reiter. An architecture for data-to-text systems. In Procs. of 11th European Workshop on Natural Language Generation (ENLG-2007): pp.97-104.

Rubinstein 1998. Ariel Rubinstein. *Modeling Bounded Rationality*. MIT Press, Cambridge Mass.

Rubinstein 2000. Ariel Rubinstein. *Economics and Language: Five Essays*. Cambridge University Press. Cambridge.

Theune et al. 2001. M.Theune, E.Klabbers, J.R. de Pijper and E.Krahmer. From data to speech a general approach. *Natural Language Engineering* **7** (1): 47-86.

Turner et al. 2008. R.Turner, S.Sripada, E.Reiter and I.P.Davy. Using spatial reference frames to generate grounded textual summaries of georeferenced data. In Proceedings of INLG-2008. Salt Fork, Ohio, USA.

Van Deemter 2004. Kees van Deemter. Towards a probabilistic version of bidirectional OT syntax and semantics. *Journal of Semantics* **21** (3).

Van Deemter 2006. Kees van Deemter. Generating referring expressions that involve gradable properties. *Computational Linguistics* **32** (2).

Van Deemter (in preparation). Kees van Deemter. *Not Exactly: in Praise of Vagueness*. To appear with Oxford University Press.

Van Rooij 2003. Robert van Rooij. Being polite is a handicap: towards a game theoretical analysis of polite linguistic behavior. In Procs. of Theoretical Aspects of Rationality and Knowledge (TARK-9), Bloomington, Indiana.

Veltman 2002. Frank Veltman. Het verschil tussen 'vaag' en 'niet precies'. (The difference between 'vague' and 'not precise'.) Inaugural lecture. Vossiuspers, University of Amsterdam.