Signalling Games: Evolutionary Convergence on Optimality[1]

Tom Lentz
Utrecht Institute of Linguistics OTS, Utrecht University

Reinhard Blutner
ILLC, University of Amsterdam

Horn's *division of pragmatic labour* (Horn, 1984) is a universal property of language, and amounts to the pairing of simple meanings to simple forms, and deviant meanings to complex forms. This division makes sense, but a community of language users that do not know it makes sense will still develop it after a while, because it gives optimal communication at minimal costs. This property of the division of pragmatic labour is shown by formalising it and applying it to a simple form of *signalling games*, which allows computer simulations to corroborate intuitions. The division of pragmatic labour is a stable communicative strategy that a population of communicating agents will converge on, and it cannot be replaced by alternative strategies once it is in place.

1 Introduction: philosophy and empiricism

If philosophy is the justification of knowledge, one of the subjects that may be justified is empiricism as a source of knowledge. But, reversely, can empiricism justify philosophical principles? Our research is based on simulation, a form of empiricism, to test the hitherto unproven but plausible evolutionary origin of a theory of language philosophy, namely Horn's *division of pragmatic labour*.

An important aspect of linguistic theory should be the possibility to account for the origin and development of language. Some language universals might be

[1] We thank Erik Borra, Arnold Obdeijn, Jasper Uijlings, and Reinier Zevenhuijzen who contributed to a very early version of this paper written in Dutch, and to the actual computer simulations on which this article is based. The first author is supported by a grant from the Netherlands Organisation for Scientific Research (NWO), project grant 277-70-001 to René Kager.

explained by, or reduced to, properties of the brain, whether these are general cognitive properties or specific to language. As the brain can be observed, theories positing properties of the brain are, at least in theory, empirically testable.

Historical experimentation or observation, on the other hand, is virtually impossible, making it very hard to test theories on language universals and their origins. This paper, however, uses simulations of language development to overcome this problem. The universal property that the paper focuses on is a pragmatic property of language, Horn's *division of pragmatic labour* (Horn, 1984) which will be explained below. This phenomenon is described as a property that is observed universally in language use; however, it is posited philosophically as a basic principle that is used to describe natural language semantics and pragmatics. This paper attempts to add more arguments to the philosophical side of the principle by explaining its emergence as virtually inevitable under reasonable assumptions on language evolution.

It is assumed that language came into being by linking signals to meanings and vice versa[2]. Given this assumption, this article is to show that the division of pragmatic labour follows from repeated acts of linguistic communication. This is important, as a population without language cannot agree on how to develop a language. This makes it undesirable to attribute pragmatic preferences of language use to individual preferences in language users, as individuals have no reliable information on the preferences of other language users a priori. Computer simulations showed that it is not necessary to assume individual preferences to be biased towards optimal pragmatic solutions; the only 'bias' should be that effective communication is preferred over ineffective communication, but this 'bias' cannot be seen as a property that determines the strategies of individual language users, as it surpasses the level of the individual. Still, an optimal solution emerges that happens to conform to the division of pragmatic labour.

In this paper, an initial stage of unprincipled form-meaning pairings is assumed. An individual may produce a specific signal (noise) in a specific situation. This signal is his expression of that situation. However, this "language" is restricted to the level of the individual and it is quite removed from a shared communicative device.

In the next section, we explain what is meant by a Horn strategy of form-meaning pairings. Section 3 introduces Lewis' idea of a signalling game (Lewis, 1969). Subsequently, in section 4 we explain our implementation of a simulation experiment of signalling games, and in section 5 and 6 we present our results

---

[2] In the modern constructionist literature, such form-meaning pairs are called constructions (see Goldberg, 1995; Tomasello, 2003).

demonstrating how Horn's division of pragmatic labour emerges from an evolutionary mechanism in language use. Further, we discuss the evolutionary stability of certain form-meaning pairs. Section 7, finally provides a general discussion including an outlook of how the present research could be continued.

2 Horn strategy

People tend to use the simplest signals for the most common messages and more complex signals for more unusual messages. This can be seen in example (1) , below.

(1) a. John sings a song.
    b. John produces noises resembling a song.

Sentence (1a) caters for a normal situation where someone sings a song. In the second sentence something strange seems to happen. Indeed, why use a strange sentence for a message that may be catered for with a normal sentence?
This observation is expressed by the following rule: *Simple messages express normal situations and more complex messages express strange situations*. This rule was posited by Horn (1984) and is known as *Horn's division of pragmatic labour* or as *Horn's rule* or (here) as the *Horn strategy*.

Horn justifies his rule on empirical grounds, by observation. The lack of observations on its development, however, makes it hard to explain its origin. As theories gain empiric backing by a multitude of observations complying with it, we will continue this paper with a formalisation of Horn's rule that allows for observations on its development. Does Horn's rule always apply to the development of language or is it an accidental feature of the languages that we happen to observe? To answer this question, populations developing a language were simulated. In the simulations, language users (agents) interact to convey meaning in so-called signalling games. The language users are assumed to be defined by a genetic make-up, which can change over generations.

3 Signalling Games

The concept of *signalling games* (Lewis, 1969) can be summarised as follows: a population of agents communicate to each other; if they manage to interpret a message correctly, they score in the game. The signalling games paradigm is well-defined and therefore it can be put to use to simulate language development in a straight-forward way, with the addition of an evolutionary perspective.

Methodologically, the development of an evolutionary perspective can proceed in two different ways. First, there is the purely theoretical approach, by using the concept of an evolutionary stable strategy (Smith, 1982). For several interesting results that were found in this way, we refer the reader to van Rooy (2004) and Benz, Jäger, and van Rooy (2005). Second, there is the construction of explicit dynamic models of the process by which the proportions of various strategies in a population change. This approach was pioneered by Luc Steels (e.g. Steels, 1998; Steels & Belpaeme, 2004).

In this paper, we follow the second approach and we will develop a genetic algorithm that implements a signalling game, adding (to the general model) the idea that well communicating agents score points and are thus more likely to procreate.

In the model, there are two roles for every agent: sender and receiver. A sender sends a message covering the meaning he wants to transmit. The receiver interprets the message; he attributes a meaning to it. Prerequisite for good communication is that the meaning is the same in both cases; only then agents understand each other. Note that no a priori form-meaning relation is imposed. This is important, as even though one might accept the emergence of simple form-meaning pairs that have iconic value (like the imitation of an animal’s call to signify that animal), assuming iconic "words" is in no trivial way sufficient to explain full-fledged languages. In addition, it also fails to explain Horn's division of pragmatic labour as the connection between simple form and common meaning is not iconic. To the human observer, that relation might be "logical", but that can be explained wholly by the fact that this is what we observe and/or have acquired, and therefore begs the question.

It is important to remark that communicating and procreating agents should not be seen as models of actual humans; the agents are way to simple and there is no evidence as of yet to indicate that pragmatic strategies are encoded directly in the genome; without any biological underpinning, such an assumption is therefore far too speculative. The simulations are meant to illustrate the high likelihood that pragmatic strategies, when adapted to communicative needs, converge on Horn-like states, thereby creating both a shared language without prior conference and without individual properties.

4 Formalisation of signalling games

4.1 Evolving agents in the signalling game

In the simulations, agents obtain points for each message they rightly transmit or interpret. Those scoring most points are most likely to survive and procreate. An agent is fully defined by his communication strategy. His only task is to communicate and that is the sole thing that matters for his survival and procreation. The game is bilateral; each agent both sends and receives.

Offspring of two agents will have a communication strategy combining both strategies. This might be the origin of a common language, but that is merely accidental and only holds for children with similar parents, not in general. Obviously, this simplification of procreation does not model human children of which the parents speak a different language; a child raised in a bilingual situation is most likely to speak both languages. The offspring of agents with different strategies combine the strategies of their parents; one could also think of this as a probabilistic simplification, as the strategies are independent of each other.

4.2 Domain of communicative interaction

The game itself is a simplification of human communication. Agents may find themselves in two different situations, the *normal situation* and the *deviant situation*. In the model, this amounts to a desire of the agent to express the *normal meaning* and the *deviant meaning*. Agents may transmit two different signals, the *simple signal* and the *complex signal*. The names for the situations and for the meanings are mere labels; they have no properties that are different, the only difference is that they are not the same instance of the same class.

Each combination of meaning and signal is allowed. This leads to 2 (meanings) x 2 (signals) = 4 possible sending strategies as well as 4 receiving strategies. Together this leads to 4 (sending strategies) x 4 (receiving strategies) = 16 communication strategies.

The game is played as follows: two agents meet and communicate. One starts to speak. First his situation is normal, so he communicates the normal meaning in the form relevant to his strategy. The second agent interprets this signal with the meaning relevant to his strategy. If indeed this is the intended, normal meaning, both score one point. Next, the sender communicates the deviant meaning and the same happens. Subsequently, the roles are inverted. In this way both agents can score up to four points per game.

4.3 Representation of the agents

An agent is represented by a bitstring (a series of bits, i.e. zeros and ones). This bitstring represents his communication strategy. The bitstring of an agent with the Horn strategy, for instance, is 0101. The bits are 0 for simple and for normal and 1 for complex and for deviant and the positions indicate the following:

1. the signal used for the normal meaning (simple, 0, or complex ,1)

2. the signal used for the deviant meaning (simple, 0, or complex, 1)
3. the meaning attributed to the simple signal (normal, 0, or deviant, 1)
4. the meaning attributed to the complex signal (normal, 0, or deviant, 1)

These bitstrings may be considered to be the agent's genome. The bitstring can be transformed in a more insightful graph, as shown in Figure 1.

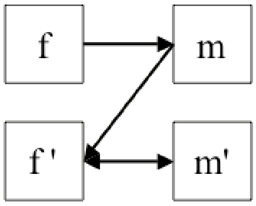

Figure 1: Example of an agent type. The arrows indicate the connections from forms to interpretations and from interpretations to forms, respectively

The total set of sixteen stategies is shown in Figure 2 in a schematic form.

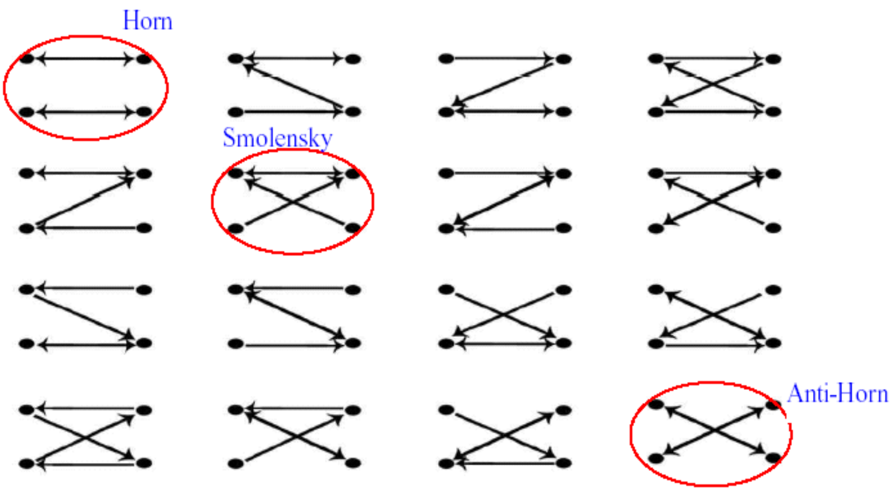

Figure 2: All communication strategies (agent types) in a schematic form. Three strategies are of special interest: (a) the Horn strategy as explained in section 3, (b) the Smolensky strategy reflecting the initial state of a learner (everything is assumed to be simple), (c) the anti-Horn strategy, which can be seen as the complement of the Horn strategy.

4.4 The selection mechanism

A selection procedure always chooses agents who survive and agents who procreate. The probability to be selected is based on the agent's number of points, it cannot be excluded that an agent with few points survives and procreates, though the chances are slim.

The number of points scored by two agents *a* and *b* is determined by adding up the two meanings *i* (i.e. the normal and deviant ones), as shown below:

$$\text{Points(a)} = \text{Points(b)} = \Sigma_i\ \delta[\ H_b(S_a(i))\ ,\ i\ ] + \delta[\ H_a(S_b(i))\ ,\ i\ ],$$

where $$\delta(x,y) = \begin{cases} 1 \text{ if } x = y \\ 0 \text{ if } x \neq y \end{cases}$$

$H_x(\,f\,)$ gives the interpretation by agent x for a signal *f*.
$S_x(\,m\,)$ gives the signal used for it by agent x for a meaning *m*.

The following mechanism of *procreation* has been used: if two agents are selected as having offspring, a random point in the bitstring is chosen. Two new agents, children of two parents, are created. For the one child, the bits to the left of the point originate from the first parent, the bits to the right from the second parent. For the other child, the bits to the right of the point originate from the first parent, the bits to the left from the second parent.

At birth, agents might undergo a *mutation*. Every bit has a tiny chance to mutate after having been determined by the characteristics of its parents. This probability, the *mutation ratio*, was set to be 0.01 by default.

Convergence of a population is related to a strategy being predominant within a population, as it tends to be advantageous for agents to use that strategy, since they tend to be well understood and they do understand well. The same holds for humans: in England it is useful to speak English, because many people will understand it; in Japan it is better to speak Japanese. The utility of a language is not a function of the language alone, but of the population and the language in interaction.

Agents using the predominant strategy will have more offspring, making the population converge towards that strategy. However, not all strategies make communication optimal in a homogeneous population; in fact only two strategies allow to get the maximum of 400 points, namely the Horn and the anti-Horn strategy. Horn-agents are those agents that communicate the normal meaning with the simple form, the deviant meaning with the complex form, and *vice verse* for interpretation (see Figure 2). Anti-Horn is the opposite.

As convergence does not necessary imply that a population is communicating in an optimal way, convergence and stability are defined as two separate concepts. Convergence is characterized as follows:

> *Convergence of a population takes place if a considerable majority of that population shares the same strategy in the limit.*

The percentage constituting a considerable majority may be adjusted. Stability is informally characterized as follows:

> *A strategy is stable if an already existing majority of users of that strategy cannot be overwhelmed by another strategy, i.e. if the majority uses a strategy that is the best given the population.*

This means that convergence towards a stable strategy is irreversible. In our evolutionary game a stable strategy may be compared to a Nash equilibrium. The Nash equilibrium is a concept from game theory developed by the mathematician John Nash (Nash, 1950). Two players find themselves in a Nash equilibrium if none of them gains by changing his behaviour. In our game it is not the interaction between two agents that matters, but the interaction between all agents. Moreover, the players themselves cannot change their behaviour. That means that in a way the entire population finds itself in equilibrium, a stable situation, with a certain strategy, if no strategy exists allowing an individual to score more points. This individual could come into existence by mutation or procreation; this is not likely to happen, because the mutation rate is low and in case a population is homogenous, children will usually be copies of their (identical) parents. However, if a different strategy is more successful, it can overtake the population given enough strategies, as it is more likely to procreate. In case of a Nash equilibrium, none of the fifteen possible alternative strategies is more successful in a population using the prevailing strategy.[3]

Finally, we should note that a stable strategy in the simulations amounts to a shared language; if a population converges to it, the language is shared.

[3] For readers interested in precise definitions, we refer to Weibull (1995), van Rooy (2004), and Benz et al. (2005).

5 Simulations of signalling games strategy evolution

Given the above formalisations, a series of computer simulation was run to assess evolution and emergence of the signalling strategies. The simulations started with a start population of 100 agents. All agents interact (play the signalling game) with all other agents, and acquire points for successful communication, as described above. The selection mechanism is then applied to yield a new generation of agents; 85% of the agents perish, and are replaced by new agents, that are children of existing agents (cross-overs), with a 1% chance of mutation. All simulations were iterated 100 times, to assure findings were not due to chance.

5.1 Stability of Horn and anti-Horn

Procedure
In the first group of simulations, all agents had used the Horn strategy at the start of the evolution. Apart from that, the settings as described above are used.

Results
The Horn strategy is an excellent strategy throughout the evolution. The fitness of agents is close to 80 percent of the maximum (averaged over the 100 simulation experiments). The fitness percentage stays close to the maximum value, as most agents use the Horn strategy, even though mutation adds a low percentage of deviant strategies every generation. The population always converged towards Horn, in all 100 evolution simulations. The anti-Horn strategy never reaches dominance.

The mirror imagee emerges when the initial population is anti-Horn; in that case Horn never emerges as dominant strategy, and the population converges to anti-Horn in all of the 100 simulated evolutions.

Discussion
The fact that the Horn and anti-Horn strategies dominate populations that were initially already Horn respectively anti-Horn is not utterly surprising, but it is interesting to see that the result is so persistent. In all evolutions, Horn stays Horn and anti-Horn stays anti-Horn. This shows that the Nash equilibriums that these two strategies exhibit, are very relevant to the evolution of the communication strategy. It can be explained why only these two strategies are stable. Communication between two Horn agents is the best possible and thus scores the maximum number of points. The same goes for two anti-Horn agents. All other strategies show weaknesses of fitness. One of the weaknesses is that

the send-strategy does not distinguish between normal and deviant meaning, if the send strategy is: “use the simple form for the normal meaning”, but at the same time: “use the simple form for the deviant meaning”. This strategy does not distinguish between situations. An agent using this strategy will never be understood in the best possible way, because there is no way of telling what meaning it tries to signal. Another weakness is when the send and receive strategies do not match. The send strategy is: “use the complex form for the normal meaning”. The receive strategy is “interpret the complex form as the deviant meaning”. Agents using such a strategy cannot communicate in the best possible way with identical others; if they want to communicate the normal meaning, the other agent will understand it as being deviant. However, this weakness only arises in. By chance, one agent with this strategy could perform very well, or even perfectly, if his strategies happen to go well with those of the other agents. However, the successful agent will have more offspring, filling the population with more and more agents with identical strategies; these agents cannot communicate very well. This effectively caps the total percentage of agents with inconsistent strategies; the cap is far below any reasonable convergence threshold, i.e., below 50%.

5.2 Starting simple or at random

Tesar and Smolensky (1998, 2000) describe a population of agents all using the simple strategy. In their terms, markedness is initial more important than faithfulness; difficult things are avoided. In the signalling game described above, it means that a simple signal is used for all meanings, and every signal is interpreted as a normal situation, a simple meaning. The idea is that a population will start with such a strategy, possibly because complex signals did not exist in former generations, and deviant meanings could not be distinguished from simple situations. If such a population would converge towards the Horn strategy, model and actual language would nicely match. Two series of 100 simulations were done to test the influence of the start population.

Procedure

To see whether the Nash-equilibriums strategies are indeed dominant, another series of hundred evolutions was simulated, now starting with the simplest strategy, the Smolensky strategy. That strategy is based on Smolensky’s idea that in acquiring language, simple forms and interpretations are preferred. Taking that to mean that no meaning is expressed with the complex form and no form is interpreted to be deviant, this conforms to the strategy highlighted in Figure 2 as Smolensky strategy. The whole procedure was repeated with mixed strategies, in which agents in the start population have a random strategy, chosen from all sixteen possible strategies.

Results

In the evolutions starting with the Smolensky strategy, the Horn and anti-Horn strategies emerge, but neither of them seems to be able to tip the balance and dominate the population. As can be seen in Figure 3, either the Horn or the anti-Horn strategy manages to climb to slightly below 50%, but not further. This is usually due to the opposite strategy, or in a few cases a variant of that, that cannot improve by changing towards Horn or anti-Horn (all of these cases can one-by-one be explained, but the details are left out as they are not important for the general argument).

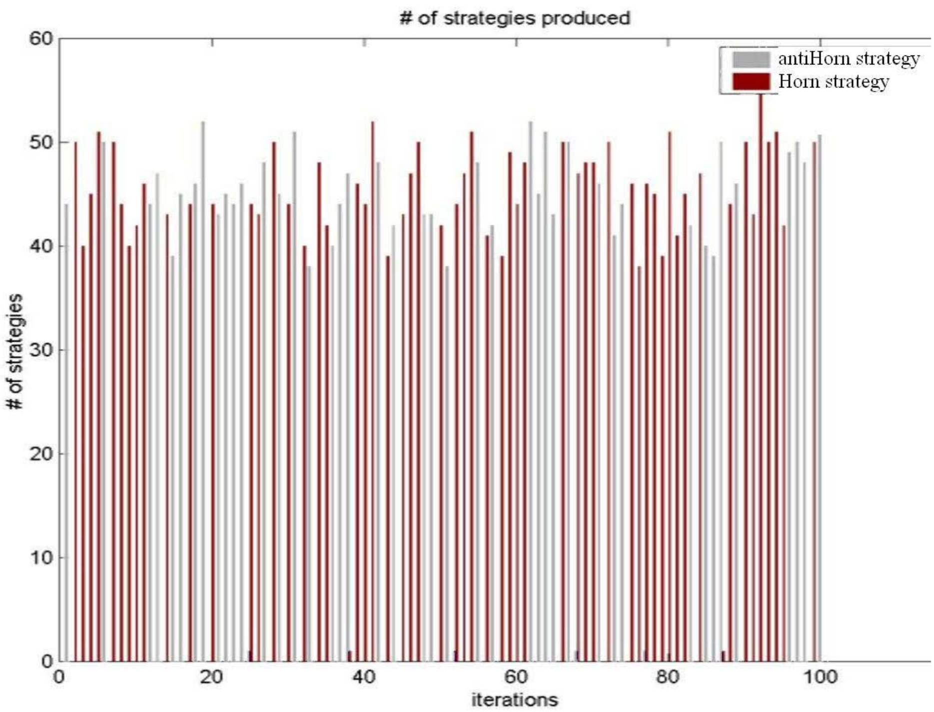

Figure 3: Percentage of agents using Horn or anti-Horn strategy at the end of 100 simulations (iterations) starting with homogenous Smolensky population.

Further, it was found that in simulations starting with mixed populations, convergence was almost perfectly divided amongst Horn and anti-Horn strategies.

Discussion
The fact that the Horn and anti-Horn strategies always emerge shows their evolutionary force. The fact that none of them is able to dominate a population shows that both are very stable, and could be compared to a population that is divided into two groups. No agent can communicate outside its group, but if its offspring changes towards the other group, it looses as much as it gains. Note that agents that do not communicate could still have offspring together, as procreation only depends on individual fitness scores. (It is up to the reader to assess if that is in accordance with real evolution.)

It is likely that mutations make some agents have offspring in the other group, but as this happens in both groups, neither of the groups is able to take advantage of that fact. This simulation, together with the above simulations, shows that evolutionary stability does not enforce an outcome, but that evolution can only go towards the two Nash-equilibriums.

6 The difference between Horn and anti-Horn

In the simulations described above, convergence towards both Horn and anti-Horn occurred (in a 50:50 ratio), whatever the start population was (except that Horn does not converge to anti-Horn, nor *vice versa*).

Of course, just by sheer definition, the Horn strategy could never dominate the anti-Horn in the simulations as described above; the meanings and the forms are not different in any way, and therefore both strategies are equivalent. However, by introducing conditions that follow from the definition of deviant and marked, the population can be made to converge towards the Horn strategy and not towards the anti-Horn strategy. These conditions are:

1. the use of the complex form is costly, in terms of points, and
2. the deviant meaning occurs less frequently than the normal meaning.

Thus the population is stimulated to use the simple form more frequently than the complex one. Since the normal meaning occurs more often than the deviant one it may be expected that the simple form will be linked to the normal meaning and the complex form to the deviant meaning: the characteristics of the Horn strategy. The conditions comply with Horn's description of "pragmatic labour": as little effort as possible will be made. However, this is less trivial than it seems. The use of the most economic strategy hardly makes any sense if one is not understood. Therefore the cost and benefits have to be balanced in some way. The assertion that frequency of the situations is non-identical, is not a complication of the original assumptions; it only formalises the distinction that was already put forward in the original definition of Horn's division of pragmatic labour.

6.1 Adding costs to the model

The method to attribute points is somewhat extended:

Points(a) = Points(b) =
$\Sigma_i P(i)\ [\delta(H_b(S_a(i)), i) - k(S_a(i)) + \delta(H_a(S_b(i)), i) - k(S_b(i))]$
P(i) is the probability of meaning/situation i;   k(f) gives a cost for signal f.

For the simple meaning we assume a probability of 1, the probability of the deviant meaning was varied between 0 and 1. In the same way the cost of the simple signal was assumed to be 0, and the cost of the complex signal between 0 and 1 (to prevent that a successful conversation would produce a negative yield).

6.2 Simulations with costs and probability added

Procedure
The procedure was the same as in the simulations of section 5, but with the cost function as just described. The cost differences are relatively low cost differences (0.8 for simple versus 1 for complex), as are the differences in the probability of normal and deviant meaning (0.8 for normal versus 1 for deviant). All four different simulations were repeated with cost and probability.

Results
When starting with the Smolensky strategy, the population eventually converges on the Horn strategy (although the number of evolutions needed for converges can be high when starting with the anti-Horn strategy). The anti-Horn strategy stops to be stable and a population of agents using the Horn strategy comes into being in 98% of the cases.

The simulations with mixed strategies end similarly. This also holds for start populations of Horn only (that do not change). It does not hold for the anti-Horn strategy under the present settings.[4]

[4] It is possible to force the Horn-strategy to emerge from an anti-Horn population with a combination of extreme cost and probability differences and an extremely high mutation rate. This is disregarded, as it violates the assumptions of the modelling realistic evolution.

7 General conclusion

The simulations show that without cost the evolutionary system either converges towards the Horn strategy or towards the anti-Horn strategy. These are the only two strategies with which the users communicate in the best possible way; they are also Nash equilibriums.

A simple addition to our model of communicating agents, namely costs and probabilities, is enough to show why and how the division of pragmatic labour makes sense. The connection between improbable situations and more involved (costly) signals emerges from simple interactions between agents that individually do not decide on the division of pragmatic labour.

The simulations with cost show how the population can converge towards the Horn strategy. This strategy is more efficient than the anti-Horn strategy and costs less. It is a combination of the strategies “minimise cost” and “maximise utility”.

The results we obtained are not extremely surprising. It is easy to see that the Horn strategy yields most points under the given circumstances. However, the conclusion that evolution tends to go into the direction most favourable to the entire population is not trivial; whether the Horn strategy is indeed the optimal solution remains to be seen for every agent. In addition, as the prisoner’s dilemma shows, without prior conference a group of agents might not converge on the strategy that is optimal for the population.

The most interesting result is, however, that the evolution model is able to abandon an already chosen direction (a strategy used by the majority of the initial population) and to end up at the Horn strategy, for 15 of the 16 possible strategies. This all happens without explicit co-ordination by the agents; moreover, the agents themselves do not weigh the possibilities.

This research may be continued in a number of ways, none of which trivially lead to the same result, even though they are likely to show similar outcomes. A more interesting and realistic model is possible by having the agents to use more different signals and to put them in more different situations. In addition, it is more realistic if an agent would be able to choose from various signals, each with a certain probability of being sent in a given situation.

The simulations presented here assume that every agent speaks with the same frequency to any other agent. It would be logical to make agents speak to others located in their neighbourhood more than with agents far away. This would lead to subgroups that understand each other well, but members of different subgroups less well, as is the case in dialects.

For philosophy’s sake it would be interesting to research the value of testing a theory in this way, since the circumstances are dictated by the theory itself. However, the theory used to be a principle that explained other semantic/pragmatic phenomena (like implicatures); the simulations and formalisations presented here show that the principle is consistent and does not need an assumption of innateness; it arises from very basic and assumedly uncontroversial formal translations of the definition itself. The non-iconic connection between simple form and normal meaning, paired with the coupling of marked form to marked meaning, emerges in evolution, and it does not have to be “designed in” anywhere in the individuals’ systems. It is both a possible outcome and the best outcome for signalling games in an evolutionary perspective.

Finally, it should be noted that similar results could be found by a paradigm called iterated learning (e.g. Kirby & Hurford, 2002) which can be seen as an alternative approach to cultural evolution. An important research objective is to adjust the existing methods of cultural evolution and to apply them to empirically investigated situations of language change. This necessitates first of all a clarification of the relationships between iterated learning and Steel’s recruitment theory (see Steels, 1998), as well as between the main internal constituents of either of them.

8 References

Benz, A., Jäger, G., & van Rooij, R. (2005). *An introduction to game theory for linguists*. Houndsmills, Basingstoke, Hampshire: Palgrave Macmillan.

Blutner, R., Borra, E., Lentz, T., Obdeijn, A., Uijlings, J. & Zevenhuijzen, R. (2002). Signalling Games: hoe evolutie optimale strategieën selecteert. In *Handelingen van de 24ste Nederlands-Vlaamse Filosofiedag*. Amsterdam: Universiteit van Amsterdam.

Goldberg, A. E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*: University Of Chicago Press.

Horn, L. (1984). Towards a new taxonomy of pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, form, and use in context: Linguistic applications* (pp. 11-42). Washington: Georgetown University Press.

Kirby, S., & Hurford, J. (2002). The Emergence of Linguistic Structure: An overview of the Iterated Learning Model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the Evolution of Language* (pp. 121-148). London: Springer Verlag.

Lewis, D. (1969). *Convention: A Philosophical Study*. Princeton: Harvard University Press.

Nash, J. F. (1950). Equilibrium points in N-person games. *Proceedings of the National Acadamy of Sciences of the United States of America 36*, 48-49.

Smith, J. M. (1982). *Evolution and the Theory of Games.* Cambridge: Cambridge University Press.

Steels, L. (1998). The origins of syntax in visually grounded robotic agents. *Artificial Intelligence 103*, 133–156.

Steels, L., & Belpaeme, T. (2004). Coordinating Perceptually Grounded Categories through Language. A Case Study for Colour. *Behavioral and Brain Sciences*.

Tesar, B., & Smolensky, P. (1998). Learnability in Optimality Theory. *Linguistic Inquiry 29*, 229-268.

Tesar, B., & Smolensky, P. (2000). *Learnability in optimality theory*. Cambridge Mass.: MIT Press.

Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, Mass.: Harvard University Press.

Van Rooy, R. (2004). Signalling games select Horn strategies. *Linguistics and Philosophy 27*, 493-527.

Weibull, J. W. (1995). Evolutionary Game Theory. Cambrige, Mass.: MIT Press.