Robustness, Optimality, and the Handicap Principle*

J. McKenzie Alexander

Department of Philosophy, Logic and Scientific Method
London School of Economics and Political Science

7 February 2009

In 1980, the British gothic rock band Bauhaus released their debut album *In the Flat Field*. Track six, entitled "Small Talk Stinks", complained about the meaninglessness of middle class conversation. Although I suspect it wasn't intentional, in that song Peter Murphy identified a curious phenomenon: the human ability to communicate vastly exceeds that of any other species on the planet, and most of the time we seem to waste it. Although many species communicate, ranging from simple bacteria to our nearest evolutionary ancestors, the higher primates, only humans have developed the ability to communicate to the extent that we do. Yet when we reflect upon the content of our conversations, many of them, ultimately, seem to be about not much of anything.

Surprisingly, the fact that so much of our conversation seems to be about not much of anything is, according to Dessalles (2007), the answer to the question of *why* language evolved:

> If evolution endowed us with language and the cognitive means associated with it, it was not for the purpose of speculating about the world into which we have been brought, of collaborating on the building of bridges or rockets or even devising systems of mathematics. It was so that we could chat. (Dessalles, 2007, pg. 269)

The last sentence must be taken quite literally: evolution endowed humans with the ability to speak not because of the benefits conferred by exchanging information, or theorizing, or constructing explanations of events, but so that we could *primarily* engage in idle banter (for signalling and social bonding purposes). This is a startling, and counterintuitive, claim.

*Forthcoming in *Biology and Philosophy*.

Of course, saying that something is counterintuitive hardly constitutes an argument against it: many of our intuitions are misleading if not flat-out wrong. (Think of how much effort it takes to stop thinking of the world in Aristotelean terms.) But a counterintuitive claim compels us to take a very close look at the reasons given in support of it. I shan't go through all of the reasons why Dessalles thinks that chatting is the ultimate factor lurking behind the evolution of language, but I will examine three core elements of his argument:

1. His political model of the evolution of language;
2. The claim that evolution selects for locally optimal outcomes, and that language is locally optimal for some adaptive function;
3. The claim that Zahavi's Handicap Principle poses a problem for the evolution of human language, and that the notion of conversational relevance (and the cost of being relevant) provides a solution.

In doing so, I will argue for the following:

1. That the political model of the evolution of language is not robust;
2. That evolution need not select locally optimal outcomes, and that consideration of the evolutionary dynamics of sender-receiver games gives us little reason to think that language need be locally optimal for some adaptive function;
3. That recent work on the evolution of signalling systems challenges the view that honest and reliable signals can only exist if they are costly, and hence Zahavi's Handicap Principle may not pose a problem for the evolution of human language.

The upshot is that we are left uncertain as to what extent Dessalles's account is ultimately supported. I argue that, although it may very well be true that language evolved, in part, because of the benefits brought about by chatting, that it is too hasty at this point to suggest that such benefits are the primary reason language evolved.

1. The political origins of language

Dessalles suggests that idle banter is a way for people to signal their quality as information detectors. Being good at "small talk" reflects people's ability to report on potential changes in, and salient features of, the surrounding environment. Sterling conversationalists win prestige within their group, thereby increasing their social status. However, in order to avoid the problem of people exaggerating the salience of the situation of which they speak, "[h]earers try to assess accurately the quality of the information presented to them so as to 'reward' it properly through the granting of status." (Dessalles, 2007, pg. 339). People with a higher status are considered to receive a fitness enhancement as a consequence.

Why does Dessalles call his model of the evolution of language a *political* model? The reason is that we are evolved from social primates, and naturally form coalitions for protection and existence. Which groups we choose to associate with, and whom we choose to let into our group, can be viewed as a kind of political act. Given this, "[s]ince belonging to a coalition is of such vital importance for individuals, what are the criteria on which they choose each other?" (Dessalles, 2007, pg. 347). We know, through the work of Franz de Waal and others, that chimpanzees pick members to belong to their coalitions based on physical size and strength. This makes sense because chimpanzee coalitions tend to be small so adding a single individual can make the difference between a successful or unsuccessful competition between coalitions. However, *we* probably use a different criterion, since our group sizes tend to be sufficiently large that adding a single person would be unlikely to make a considerable difference. Dessalles proposes that we consider the ability to be relevant as an important criterion for group membership.

Given that we know relatively little about the conditions under which language evolved, it would be nice if Dessalles's model was *robust*, as this would show that the particular assumptions made do not matter much.[1] This notion of robustness is taken from population biology:

> we attempt to treat the same problem with several alternative models each with different simplifications but with a common biological assumption. Then, if these models, despite their different assumptions, lead to similar results, we have what we can call a robust theorem that is relatively free of the details of the model. Hence our truth is the intersection of independent lies. (Levins, 1984, pg. 20)

Consider the following agent-based implementation of Dessalles's political model of language formation, drawn heavily from Dessalles (1999):

- Each agent a_i possesses two traits: $g_i^1 \in [0,1]$, which measures the propensity to "speak relevantly" about some issue, and $g_i^2 \in [0,1]$,

[1]Weisberg (2006) provides a nice example of robustness analysis regarding the Volterra Principle.

which measures the propensity of the agent to bestow status upon another person when they speak relevantly[2].

- Each agent a_i has an inherent base ability $q_i \in [0, 1]$ to speak relevantly.
- We begin with a population $P = \{a_1, \dots, a_N\}$. For every round of interactions except the very first, we partition P into a set of coalitions $\mathcal{C}$.
- Each person a_i in the population is paired with k other people at random, with preference given to interacting with members from her own coalition C_i (if a_i and a_j belong to the same coalition, then $C_i = C_j$).[3] Let $\eta_i^t = \{a_{i_1}^t, \dots, a_{i_k}^t\}$ denote the people with whom a_i is paired with at time t. For each member $a_j \in \eta_i^t$, the following takes place:
  1. a_i tries to speak relevantly with a_j. Agent a_i succeeds with probability g_i^1.
  2. If a_i speaks relevantly with a_j, then a_j benefits by $Gq_ig_i^1$ (where G is a global constant) and a_i incurs a cost of $C_1q_ig_i^1$.
  3. If a_i spoke relevantly, then a_j may bestow status upon a_i. He does so with probability g_j^2.
  4. If a_j rewards a_i with status, then a_i receives a status increase of $Rg_j^2q_ig_i^1$ (where R is a global constant set at the beginning of the simulation) and a_j incurs a cost of $C_2g_j^2q_ig_i^1$.
- After all interactions have taken place, a round of "coalition competitions" takes place as described in Dessalles (1999).
- Each agent appearing in a nondegenerate coalition has her status multiplied by the relative aptitude of her coalition (see H7, in Dessalles, 1999).
- After all agents in the model have had all of their interactions, we replicate agents using a Moran process.

[2] Superscripts denote indexes rather than exponents.

[3] Given the dynamics of coalition formation, it is possible that some coalitions will be degenerate (i.e., containing only one member). When a coalition is degenerate, the people with whom a_i is paired are drawn uniformly from $P \setminus \{a_i\}$. (And it is allowed for a_i to interact with the same person more than once.) If coalition C_i is not degenerate, then *each* of the k interactions for a_i is determined as follows: a_i is paired with someone drawn uniformly from $C_i \setminus \{a_i\}$ with probability 0.6, and otherwise with someone drawn uniformly from $P \setminus \{a_i\}$.

Define the *fitness* of a_i at the end of generation t to be the sum of all costs, benefits, and status increases received by the agent over the last period. Let the *status* of the agent at the end of generation t simply be the sum of all the status awarded to him by others as a result of his interactions. The status is thus a separate quantity from the fitness of the agent, but the status of the agent contributes to the agent's fitness.

A *Moran process* is often used for modeling evolutionary dynamics in finite, unstructured populations. At time t, each individual a_i is assumed to have a nonnegative fitness f_i, with at least one individual possessing a positive fitness.[4] Fitness are converted into replication probabilities through normalization: the probability that a_i will replicate is simply $\frac{f_i}{\sum_j f_j}$. One individual, say a_i, is selected at random for replication and another individual, say a_j, is selected at random, removed from the population, and replaced by a clone of a_i.[5] A Moran process has two useful properties: first, it keeps the size of the population fixed. Second, it is a useful evolutionary dynamic because one can establish analytic results concerning its long-term convergence behaviour.[6]

Figure 1 illustrates the outcome of one simulation in the agent-based political model of the evolution of language. Notice that, in contrast to the result reported in *Why We Talk* (see figure 2) we do not get convergence to a 100% "communication level". What we see, rather, is that the average value of g^1 in the population, representing the probability of people speaking relevantly in an interaction, initially declines, then slightly increases, then gradually declines again over the 6000 generations.

However, if we pick other values for the constants, as in figure 3, we find different behaviour. These alternate values may seem more reasonable as they set the cost of both speaking relevantly and bestowing status upon others rather low. In this second simulation, after some initial transient noise is driven out, the population slowly increases the mean values of g^1 and g^2 until it arrives at a point where everyone in the population is identical in type. Note, though, that little significance should be attached to this: because the evolutionary dynamics we are using do not permit the introduction of new types, a Moran process, given enough time, will eventually converge to a state

[4]Because individuals incur a cost when they attempt to speak relevantly, I assume all agents have a baseline fitness of 3. This insures that even if someone chooses to speak relevantly to others, and receives no benefit or status increase, he still has a positive fitness.

[5]Here, I shall assume that a_j is selected using a uniform probability distribution over the entire population.

[6]Pawlowitsch (2007a) models the evolution of a proto-language in a finite population using a Moran process and shows that "efficient proto-languages are the only strategies that are protected by selection." See also the discussion in section 2.1.

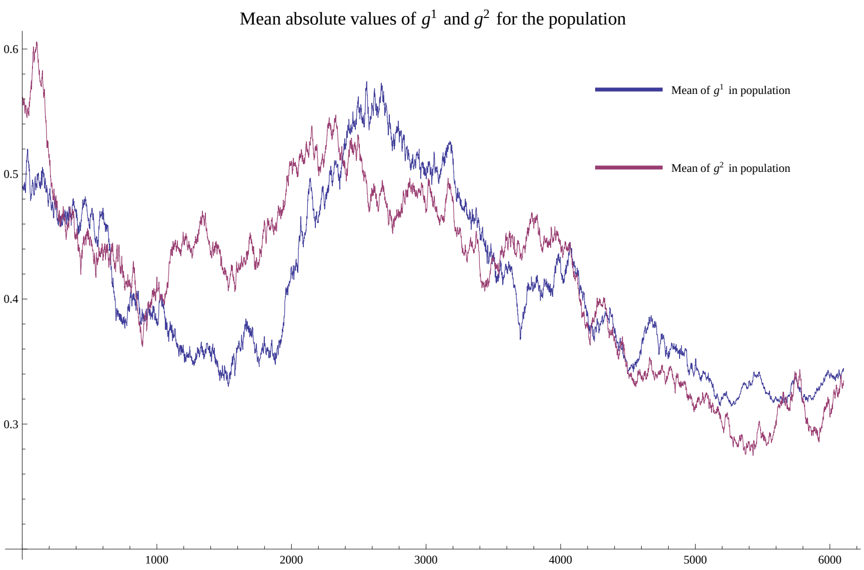

(a)

Figure 1: Time-series plot of the mean values of g^1 and g^2 for the population in the agent-based variant of Dessalles' political model for the origin of language. Population size of 150, $(G, R, C_1, C_2) = (2, 2, 1, 1)$, and three interactions initiated by each agent in each round.

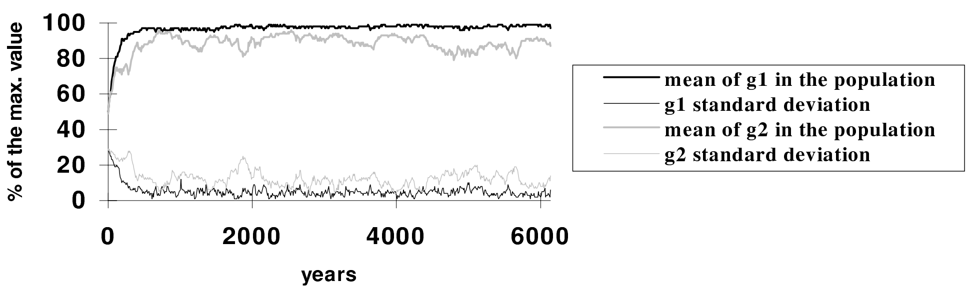

Figure 2: The emergence of language behaviour (Dessalles, 2007)

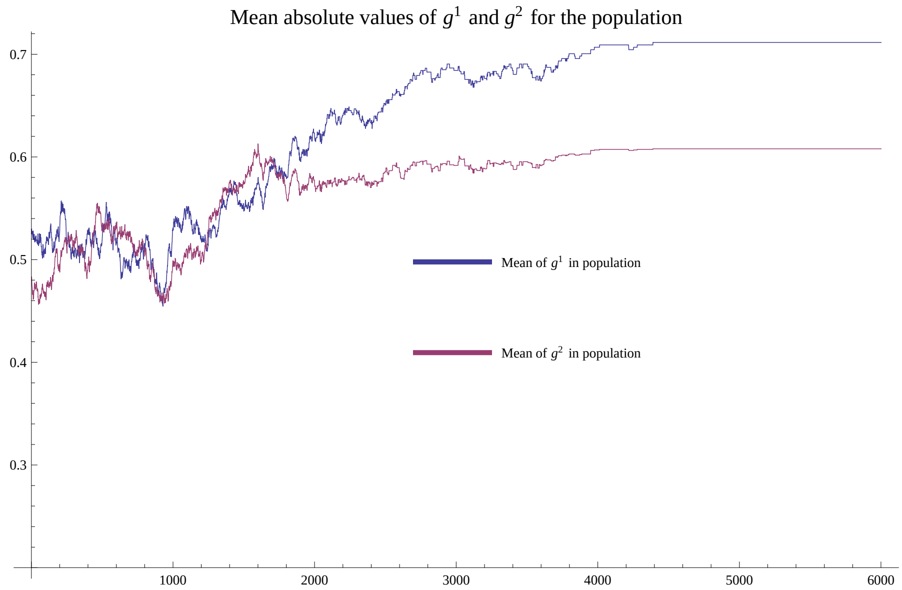

(a)

Figure 3: Time-series plot of the mean values of g^1 and g^2 for the population in the agent-based variant of Dessalles' political model for the origin of language. Population size of 150, $(G, R, C_1, C_2) = (1.5, 1.75, .5, .6)$, and three interactions initiated by each agent in each round.

where everyone is identical. For the simulation represented in figure 3, this consists of individuals have the following values: $g^1 = 0.71148$, $g^2 = 0.6079$, and $q = 0.852001$.

The result of figure 3 appears to replicate the result of Dessalles (2007, pg. 350). However, this reproduction is merely due to the random selections made in the Moran process: a second run, initialized with similar initial conditions and identical values of the key constants, produced the outcome illustrated in figure 4. Notice how, in that figure, although the mean value of g^1 does increase to a noticeably high value of 0.8 within 6000 iterations, the mean value of g^2 *declines* to a low point of 0.2. In words: although agents are increasingly inclined to "speak relevantly", not too many people care enough in order to be bothered to award them with an increase in status.

One shortcoming of the agent-based model is that it does not permit the introduction of new strategies (or player types) into the population. What happens if we make that possible? Let us implement a process of "mutation" as follows: if mutations are permitted, each agent a_i has a chance μ of being replaced by an entirely new type of agent. Values of the critical parameters g_i^1, g_i^2 and q_i are selected at random from $(0, 1)$.

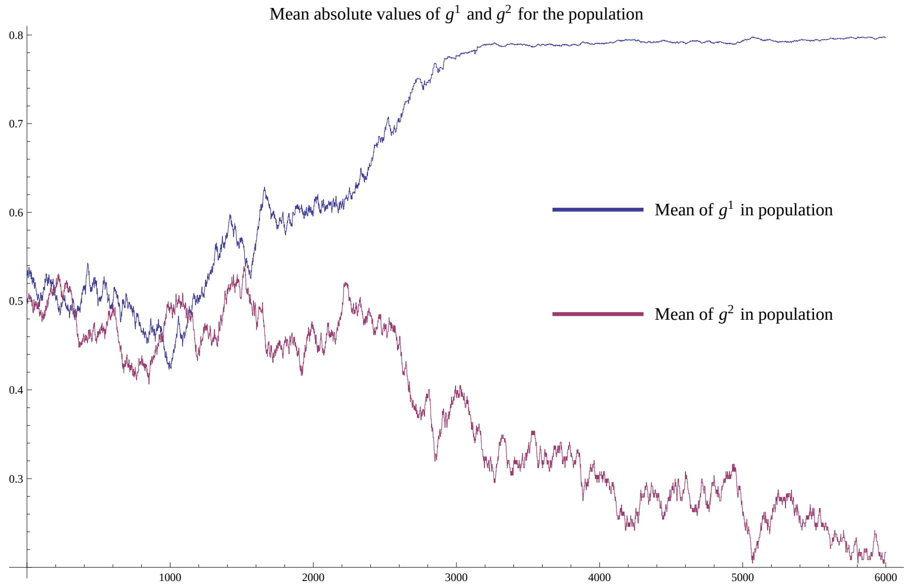

(a)

Figure 4: Time-series plot of the mean values of g^1 and g^2 for the population in the agent-based variant of Dessalles's political model for the origin of language. Population size of 150, $(G, R, C_1, C_2) = (1.5, 1.75, .5, .6)$, and three interactions initiated by each agent in each round.

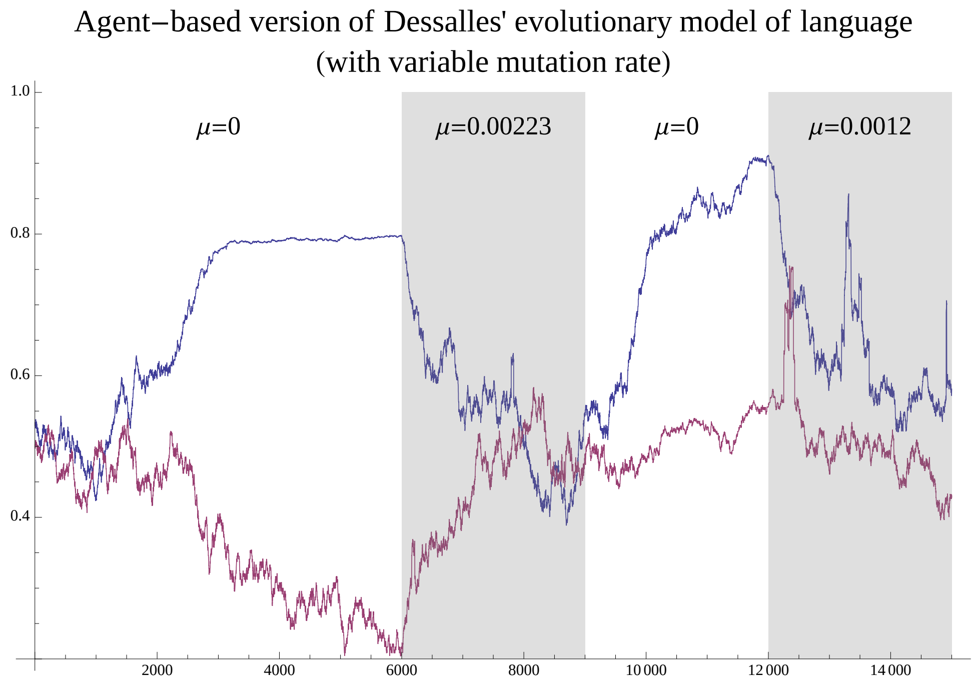

(a)

Figure 5: Time-series plot of the mean values of g^1 and g^2 for the population in the agent-based variant of Dessalles' political model for the origin of language. Population size of 150, $(G, R, C_1, C_2) = (1.5, 1.75, .5, .6)$, and three interactions initiated by each agent in each round. Variable mutation rates μ used, with values as specified in the plot.

Figure 5 illustrates how mutations serve to destabilise the apparent reproduction of the result reported by Dessalles (2007). The first 6,000 iterations in that plot coincide with that of figure 4 (it is from the same simulation). However, once the population has reached a reasonably "stable" level of nearly 0.8, for the mean value of g^1, and 0.2, for the mean value of g^2, mutations were enabled at a rate of $\mu = 0.00223$. The reason for choosing this particular mutation rate was that it ensured the appearance of approximately one mutant every three iterations[7]

Notice that the presence of even a relatively low rate of mutation serves to overrun the influence of selection: The population mean for both g^1 and g^2 is driven back into the range around 0.5 for both parameters, as shown in the first shaded region. Admittedly, if mutations are disabled again for a period of time, as done between iterations 9,001–12,000, the population mean for

[7]The population is of size 150 and $150 \cdot \mu \cdot 3 \approx 1.0035$.

g^1 (although not g^2, interestingly) increases again — to a value approaching 0.9. However, re-enabling mutations, with a rate *half* of what it was during iterations 6,001-9,000 serves to knock the population out of this state.

This shows that the outcomes of the model reported by Dessalles (2007, pg. 350) depend substantially upon the interplay between three features: (a) the particular evolutionary dynamics used, (b) the values of key constants in the model, and (c) the method by which new player types are introduced into the population. Given our ignorance of all three of these regarding the emergence of language during the early period of *homo sapiens*, this dependency of Dessalles's result upon the model should lead us to view his result with caution. Had we found that the same, or at least similar, results were produced from a variety of different evolutionary models, under a range of different conditions, we would have some confidence that we had "found the truth at the intersection of various lies". But the above results suggest that Dessalles's political model for the emergence of language may not be robust. Not only have we not found the truth at an intersection of lies, it is unclear whether we have even found an intersection.

However, even if the model *was* robust, would we be able to draw the inferences from the model which Dessalles suggests we can? I suggest the answer is no. The worry is that the political model of the evolution of language can also be interpreted as a simple model of the evolution of co-operation in a group structured context. Given how little we actually know about the origins of linguistic behavior, the political origins looks more like an interesting Just So story.

Second, the political model leaves many relevant factors unexplored. What happens if people speak relevantly in conversation (hence showing their capacity to convey information about the environment) in order to achieve membership in a coalition, but then fail to provide information to others? This possibility isn't explored in any real detail, although the general problem of language as a cooperative enterprise is discussed in chapter 16. (Dessalles suggests that Zahavi's "handicap principle" may play a role in solving this problem. One difficultly, though, is that other work in evolutionary game theory suggests that Zahavi's "handicap principle" is false as a general rule. I return to this point in section 2.2.)

Another problem with the political model is that it fails to explain the "goodness-of-fit" between the world and language. That is, how did language evolve so as to faciliate our making fine-grain distinctions between kinds of things, making true descriptive statements about the world, and issuing commands about what to do and when to do it? This is only a small fraction of what language enables, and any account of how language evolved would need to have an answer to these questions.

One natural framework for tackling this problem is one which Dessalles does not cover in his book: that of Lewis's sender-receiver games. What I will do, briefly, in the remainder of this paper is discuss some results from Lewis signalling games to illustrate how they challenge two core elements which Dessalles invokes: the notion that language is locally optimal and the Handicap principle.

2. Lewis sender-receiver games as a model of the evolution of language

Consider, as another model of the evolution of language, a two-player sender-receiver game as introduced by Lewis (1969). In such games, Nature chooses a state of the world and reveals this state to the sender, who then sends a signal to the receiver who performs an action. The outcome of the receiver's action generates a payoff for both players, the payoff depending on how well the action performed "fits in" with the state of the world.

The simplest sender-receiver game takes the form of a pure coordination game; when the Receiver matches his action with the state of the world, both he and the Sender obtain payoffs of 1; otherwise, both players receive nothing. A "signalling system", as defined by Lewis, occurs when we have an optimal matching between state and action, so that the payoff for both the sender and receiver is optimal.

Sender-receiver games provide a minimal framework for modeling how meaningless signals can acquire information. It also reveals reasons we should be concerned about assuming that language is locally optimal and satisfies the Handicap principle.

2.1. Local optimality and the evolution of language

The concept of local optimality plays a prominent role throughout *Why We Talk*. Here are a few places where it is appealed to (italics mine):

> "[L]anguage is not due to a macromutation; it serves an adaptive function for which it is *locally optimal*" (117)
>
> "Microevolution is rapid because there is open competition among individuals. In equilibrium, this competition is no longer open, as all the *best available* solutions have been found [...] Microevolutionary competition enables selection to do its work and to create a pressure which pushes evolution in a given direction, that of the next *local optimum*." (124)

> "If our species has a predisposition to use a phonological system, then the predisposition must be *locally optimal* for a biologically adaptive function." (160)
>
> Regarding protolanguage: "if protolanguage was one of the characteristic behaviours of a species of hominids, it must be possible to show that it was *locally optimal* [...], that is to say that no minor variation in the competence could have made it any better at fulfilling its biological function." (172) And also: "A conclusion that appears natural is that protolanguage is *locally optimal* for communicating meanings of a certain sort and that word order is chosen so as to facilitate the hearer's construction of meaning." (172)

In these passages, notice how Dessalles shifts in the type of claim he makes regarding the local optimality of language. The first quote, in which it is said that language "serves an adaptive function for which it is locally optimal," is making a straightforward empirical claim. Some of our evolved traits *do*, as a matter of fact, serve adaptive functions for which they are locally optimal, and our linguistic ability might be one of them. However, contrast this with the last quote where he states, "it must be possible to show that [protolanguage] was locally optimal." This isn't an empirical claim, and it is difficult to understand just what Dessalles means. Why *must* it be possible to show that protolanguage was locally optimal?

Is it possible that Dessalles is using "locally optimal" in some special proprietary sense? I suppose so, but it is hard to square that possibility with the comment he makes on page 172 regarding local optimality: "that is to say that no minor variation in the competence could have made it any better at fulfilling its biological function." Normally we say that a trait is locally optimal when it is at a local maximum of the fitness landscape. The explanatory remark that Dessalles provides is perfectly compatible with the ordinary sense of local optimality.

The reason why this is a worry, of course, that it is generally incorrect to assume the outcomes of evolution are locally optimal. While evolution may produce traits, structures, or behaviours which are locally optimal, it need not. There are at least three reasons for this. The first (as Dessalles recognizes in his discussion of Gould and Lewontin's criticism of adaptationism) is that some traits, structures, or behaviours are not selected *for* at all, but are rather evolutionary spandrals. If a spandral is locally optimal, that is a happy accident rather than an explicit product of evolutionary design.

Second, developmental lock-in might preclude the possibility of obtaining locally optimal outcomes. Brute facts about how individuals of a species develop from an embryo to an adult organism may well rule out that species settling upon the locally optimal solution to a particular adaptive problem because it is a developmental impossibility. (If one responds that the concept of "local optimality" takes into account these kinds of constraints, then the concept of local optimality, becomes tantamount to saying that anything which evolves is locally optimal *by definition.*)

Third, genetic interactions *within* individuals might prevent locally optimal outcomes from being selected for. This last point is important to appreciate because many character traits and behaviours in humans and other species arise through the interaction of multiple genes. Since each of these genes may contribute some part to several nonoverlapping biological functionings, each gene will find itself subject to selection pressure from different sources. The result from the interaction of all of these genes may then not be a locally optimal solution to any particular adaptive problem, but rather a compromise between multiple forces of selection pulling in different directions.

More importantly, though, examination of the dynamics of sender-receiver games reveals instances where the outcome of an evolutionary dynamic may *not* produce a locally optimal outcome. For example, Skyrms (2008) shows that equilibria exist in sender-receiver games which are locally suboptimal, and Pawlowitsch (2007b) proves that, for the replicator dynamics, the population may become trapped in one of these suboptimal equilibria. If the evolutionary outcome of sender-receiver games says something about the evolution of language, even in a very primitive form, we then have reason for doubting Dessalles' claims that protolanguage both *is* and *must be* locally optimal.

2.2. Zahavi's Handicap Principle and the evolution of language

Finally, consider Zahavi's Handicap Principle and the role it plays in the evolution of language. Dessalles discusses this in chapter 16 of *Why We Talk*. The problem is that:

> The first effect of speech is that it enables hearers to benefit from this information and the knowledge possessed and conveyed by the speaker. If this behaviour represented mere gratuitous assistance, it should die out rapidly through the workings of natural selection. If it represented self-interested assistance, where is the quid pro quo? (Dessalles, 2007, pg. 314)

Why should human communication not exist, according to the laws of evolution? An entrenched view in biology is that reliable signals must be costly to send. Yet communication amongst humans is effectively costless, so how could it have evolved?

Dessalles agrees with Zahavi that "the only signals natural selection can favour are the reliable ones" (Dessalles, 2007, pg. 331)[8] yet denies that it is easy to lie with words. Why? It all has to do with the purported conditions under which language evolved. In his political model of the evolution of language, speakers endeavour to speak relevantly to others, which is not easy. Dessalles points out that "hearers test the logical consistency of what they are told" (Dessalles, 2007, pg. 331) in attempts to discover shortcomings. In addition, a hearer can appeal to "trivialization" when a speaker seems to be over-egging the salience of a situation.

Thus Dessalles accepts the basic idea of the Handicap Principle, that signals need to be costly in order to be reliable. Because human communication is effectively costless, he needs to locate the cost of human communication elsewhere than in the manufacturing of the signal. The cost of communicating, then, can be found in the effort required to speak relevantly. As he states:

> According to Zahavi's general idea, communication has to be a costly exercise for speakers, if the benefits of it accrue to them. [...] The cost of behaviours of inquisitiveness and exploration, of whatever intensity, can be understood in part if we see them as a way for individuals to cull information. (Dessalles, 2007, pg. 331)

A main motivating factor for Dessalles's account of the evolution of language — putting efforts to speak relevantly at the heart of the process — might stem from taking the Handicap Principle as a general truth regarding the evolution of signalling systems.

Yet is the Handicap Principle true, in general? Recent work suggests that the story is more intricate and complex than previously thought. In an interesting paper, Hurd (1995) shows how, in a basic signal-response game, signal costs can be reduced to zero without interfering with communication. How can this happen? It has to do with the structure of the game:

> It is the discrete nature of this game that allows us to separate the handicap from the stabilizing cost. The cost-free signalling described here is unlikely to be found in situations where states and signals are continuous. (Hurd, 1995, pg. 221)

[8]However, recall our discussion in section 2.1 about sender-receiver games with pooling and partial pooling equilibria. Partial pooling equilibria are not entirely reliable because they conflate states of the world, yet are still capable of being produced by evolutionary dynamics.

Similar results have been shown to hold for other types of signalling games. For example, Bergstrom and Lachmann (1998) demonstrate that the Sir Philip Sidney game also allows honest, cost-free signals to develop under fairly general conditions. We may not need to work quite so hard to try to resolve the apparent tension between the apparently costless nature of human communication and the handicap principle.

There are other reasons to downplay the centrality of the Handicap Principle in Dessalles' account of the evolution of language. (Bergstrom and Lachmann, 1997) show that "despite the benefits associated with honest information transfer, the costs incurred in a stable costly signalling system may leave all participants worse off than they would be in a system with no signalling at all." If Dessalles is committed to the claim that evolution produces locally optimal outcomes, and that human communication is costly, it seems that he would need to demonstrate that honest information transfer, in human communication, is not one of the types of costly signalling systems identified by Bergstrom and Lachmann. If it were, and if evolution produced locally optimal outcomes, evolution might produce no signalling system at all!

3. Conclusion

In this paper I have examined three elements of Dessalles's account of the evolution of language: (1) that language evolved as a consequence of the group benefits it conferred, (2) that evolution produces locally optimal outcomes and that language is locally optimal, and (3) that the Handicap Principle presents a challenge for the evolution of human language. I have argued that each of them is to some degree problematic. Although Dessalles offers a fascinating account for how language might have evolved, it remains a difficult, and largely unsolved, problem.

References

Bergstrom, Carl T. and Michael Lachmann (1997). "Signalling among relatives. I: Is costly signalling *too* costly?" *Phil. Trans. R. Soc. Lond. B* 352: 609–617.

— (1998). "Signaling among relatives. III: Talk is cheap". *Proc. Natl. Acad. Sci.* 95: 5100–5105.

Dessalles, Jean-Louis (1999). "Coalition Factor in the Evolution of Non-Kin Altruism". *Advances in Complex Systems* 2: 143–172.

— (2007). *Why We Talk: The Evolutionary Origins of Language.* Oxford University Press. French edition, *Aux origines du langage*, first published by Hermès Science Publications, Paris, 2000.

Gould, S. J. and R. C. Lewontin (1979). "The spandrals of San Marco and the Panglossian program: a critique of the adaptationist program". *Proceedings of the Royal Society of London* 205: 281–288.

Hurd, Peter L. (1995). "Communication in Discrete Action-Response Games". *Journal of Theoretical Biology* 174: 217–222.

Levins, Richard (1984). "The Strategy of Model Building in Population Biology". In Elliott Sober, ed., *Conceptual Issues in Evolutionary Biology*, The MIT Press.

Lewis, David K. (1969). *Convention: A Philosophical Study.* Harvard University Press.

Pawlowitsch, Christina (2007a). "Finite populations choose an optimal language". *Journal of Theoretical Biology* 249: 606–616.

— (2007b). "Why evolution does not always lead to an optimal signaling system". *Games and Economic Behavior* 63: 203–226.

Skyrms, Brian (2008). "Evolution of signaling systems with multiple senders and receivers". *Phil. Trans. R. Soc. Lond. B* Working paper.

Weisberg, Michael (2006). "Robustness Analysis". *Philosophy of Science* .

Zahavi, A. and A. Zahavi (1997). *The Handicap Principle.* Oxford University Press.