John Benjamins Publishing Company

Constraining context

A pragmatic account of cognitive manipulation

Didier Maillat & Steve Oswald
University of Fribourg/Vrije Universiteit, Amsterdam

1.   Introduction

Many accounts of manipulative communication stemming from a discourse analytical perspective usually tackle the conditions under which a discourse can be termed 'manipulative' by focusing their analysis on two points. One of them is a careful examination of the external contextual settings of the particular communicative event, that is, its broad social context, which includes the social relationship between speaker and audience, their respective roles and prerogatives, the status of their respective knowledge, the purpose of the event, and so on. This somewhat *contextually external* take on manipulation, which takes into consideration the inherently social nature of manipulative communication, is summarised by Van Dijk when he posits that 'it only makes sense to speak of manipulation […] when speakers or writers are manipulating others in their role as a member of a dominant collectivity' (2006: 364).[1] A second point of focus is traditionally set on the discursive and linguistic devices which can strategically be put to use by unscrupulous speakers in order to gain consent through 'illegitimate', to some extent uncooperative, means which escape the audience's awareness, notably in the (re)production of ideological systems.[2]

While the social and discursive components of manipulative communication are evidently of great interest, we would like to take an underexplored and alternative — yet compatible — direction, by specifying this *contextual* claim from a

---

1.   To do full justice to Van Dijk's own account of manipulation, we must add here that his is a *triangulated* approach, since it takes into consideration not only the social aspect of manipulation, but also its discursive and cognitive conditions. We thus use this quote only to refer to the mainstream discourse-analytical approaches which give precedence to the situational settings in which manipulation takes place.

2.   Such is the rationale of the type of research initiated in the 1970s by the East Anglia School of Critical Linguistics, inaugurated by R. Fowler, R. Hodge, G. Kress and T. Trew with their book, *Language and Control* (1979).

cognitive perspective, which is not so much concerned with the social nature of manipulation as a communicative phenomenon, but rather with the way it functions, in particular with regard to the type of cognitive processing manipulative discourse calls for on behalf of its addressees.

In communicative interactions, and in particular in the cognitive processes of interpretation, the notion of context plays a decisive role in the meanings hearers may derive from the speaker's linguistic productions. However, we will not treat context as a *given* factor here, which is imposed onto the communicative event (or in which the communicative event takes place), but rather as a set of relevant assumptions conversational participants select as they process information in order to yield meaning. In other words, we will adopt a cognitive pragmatic approach, in the vein of the relevance-theoretic account of communication (following Sperber & Wilson 1995), and suggest that manipulative communication is foremost about exploiting the inherently fallible and heuristic-based ways in which the human mind processes information.

2.   Cognitive biases and errors: Psychological grounds for manipulation

Many trends in psychological research, over the last century, have studied the errors the human mind is likely to make in specific tasks. One of these paradigms, led by Daniel Kahneman and Amos Tversky, focused particularly on judgmental errors people make when making decisions, in an attempt to shed more light on the theory of rational decision. In particular, they showed in a series of publications that people happen to rely on certain types of *heuristics* when asked to provide probability judgments, among which the representativeness and the availability heuristics (cf. Tversky & Kahneman 1974). To take an example, they showed that when asked questions of the type 'What is the probability that object A belongs to class B?' and 'What is the probability that event A originates from process B?' (1974:1124), subjects — to a great extent — answer based on the degree to which A resembles B (representativeness heuristics), which is not sufficient evidence to conclude that a causal relation obtains between A and B. Their idea is roughly that the mind does not function like a computer which systematically and consistently processes input information, but that it takes shortcuts, and often yields outputs through fallible procedures which involve risk and thus may lead to judgmental errors.[3]

---

3.   For a detailed discussion of judgmental errors, see Kahneman and Tversky's seminal paper (1974).

Though we do not intend to get into the details of this type of study here, we use this body of research as a strong empirical basis to support the idea that we humans are fallible information processors, and that we may sometimes — due to economical constraints of efficiency among other parameters — miss some information, or at least neither look for nor take into account all relevant information when processing inputs. Although Kahneman and Tversky's research was aimed at contributing to a theory of rational decision, there is *a priori* no reason to consider that such heuristics may not also apply to other rational processes, such as language processing. This is actually an idea that researchers such as Nicholas Allott seriously take on; he argues in this respect that 'reasoning involves shortcuts, many of them heuristic […]: non-algorithmic procedures which do not guarantee reaching the right answer' (Allott 2008: 252). In particular, his aim is to show that the processing of utterances may take these 'shortcuts', sometimes at the expense of consistency and logical validity, under a global constraint of efficiency. This is also a core assumption of the model we sketch out here.

An additional psychological trend, perhaps broader in scope, known as the study of *cognitive illusions* (cf. Pohl 2004), converges in this direction by showing how humans are often subject to cognitive errors, and is meant to tackle 'a plethora of phenomena showing that we deviate in our thinking, judgement, and memory from some objective and 'correct' standard' (Pohl 2004: 1). By relying on a perceptual analogy, research on cognitive illusions considers that the mechanisms involved in thinking, managing memory and providing judgments are, as in the case of optical illusions, subject to errors. As a way of illustration, let us take the case of the so-called *Moses illusion* (cf. Erickson & Mattson 1981 for the original experiment, and Park & Reder 2004 for a global discussion of this particular phenomenon).

This is an experiment where subjects are asked the following question: 'How many animals of each kind did Moses take on the Ark?' The results indicate that as much as over 2/3 of the subjects tested answer 'two', at the same time failing to notice the anomaly (i.e. that the biblical episode of the flood is about Noah, not Moses). What this shows is that some crucial, and highly relevant, information may go perfectly unnoticed and lead to a blatant judgmental error (semantically speaking there is a referential mismatch, since it makes no sense to speak of Moses in connection with the ark). Such a biased processing of information is often referred to as *shallow processing*.

The *Moses illusion* is one among many illusions that contribute to establishing on empirical grounds the mind's fallibility. However, even if the latter does to a fair extent rely on a specific linguistic formulation, it is, as most studies of cognitive illusions, meant to illustrate that the human mind, be it in decision making, perception, memory or thinking tasks, is inherently subject to biases

and errors. Our purpose is to extend this line of reasoning to matters of linguistic processing, taken as one cognitive task among others. Such an approach which turns out to be significantly compatible with a cognitive pragmatic account of communication such as Relevance Theory (Sperber & Wilson 1995, henceforth RT). Following this line of thought, we will now turn to spelling out a few considerations about pragmatic processing involved in verbal communication in order to support our proposal, namely that manipulative instances of communication may occur precisely because our cognitive processing is fallible, thus echoing Rigotti's claim according to which 'the dynamics of manipulation are very close to the dynamics of human error' (2005:69).

3.   The fallibility of information processing: Pragmatic insights

Ever since H.P. Grice's seminal work on meaning and his exploration of what he called the 'logic' of conversation (Grice 1989 [1957 1975]), pragmaticians have been concerned with a central feature of human communication, namely the idea that producing and understanding verbal stimuli involves much more than making use of a code system. Communication, in those paradigms, is envisaged as a rational manifestation of our inferential capabilities, and the point of the latter approaches is to account for the way this rationality is instantiated in verbal communication, notably by spelling out certain principles and maxims, or, if we are to reconcile all pragmatic approaches under a broader umbrella, certain communicative standards.

A core assumption of pragmatic research is the semantic underdeterminacy of meaning. By that it is considered that linguistic utterances do not transparently encode the thought which they are meant to convey, but rather that they provide clues for the addressee to adequately contextualise and inferentially derive the intended representation. In other words, communicating is not limited to the use of a shared system of symbols, but involves inferential mechanisms of intention recovery. A basic illustration of this fact is to be found in indirect speech acts, which semantically encode a different proposition than the one that is intended by the speaker.[4] An interesting implication of semantic underdeterminacy, which contemporary pragmatics builds on, draws on the fact that at least part of the success of the communicative event rests on the hearer's responsibility in deriving

---

4.   If a speaker asks 'could you hand me that spanner?' it is normally understood that she is actually *requesting* the addressee to pass her the spanner, not that she is asking whether her addressee has the physical ability to do it. What is said, in those cases, does not correspond to what is communicated.

the appropriate — intended — content, since the linguistic utterance itself is not sufficient to derive speaker meaning.

As a consequence, we take human communication to be by definition a risky endeavour, by virtue of (i) the fallibility of human inferential capabilities (as noted above), (ii) the underdeterminacy of meaning, and (iii) the fact that a significant part of the interpretive responsibility falls on the hearer, so much so that errors and misunderstandings may occur, as witnessed in everyday conversation. One crucial feature of communication which relates to this last aspect is that interpretation always takes place by means of the contextualisation of utterances. According to RT, utterances are processed against a set of contextual assumptions which the hearer mobilises, and, to a fair extent, assumptions for which *he* is responsible, even if, when communication is successful, they turn out to be compatible with what the speaker could be taken to have in mind at the time of her utterance.[5] Achieving relevance, from an interpretive point of view, is arriving at a point where the speaker's utterance has been combined with contextual assumptions so as to yield an output representation that delivers the best ratio between the cognitive effort spent and the contextual effects it produces.[6]

As such, information processing in communication is subject to uncertainty and errors, if only because the hearer may end up selecting inappropriate contextual assumptions, and also to a resource-bound efficiency constraint balancing cognitive effort and contextual effects. We claim that manipulative communication takes advantage of this, i.e. the necessarily imperfect cognitive mechanisms of information processing. In order to sketch out our model, we will focus on mechanisms of contextual selection, and address the constraints a speaker can impose through different means, notably by (mis)leading the hearer to select some (cognitively inescapable) assumptions. As we will see, this is achieved by constraining the effort/effect ratio of the selection.

Coming back to the Moses illusion example, Allott and Rubio Fernández (2002) explained the phenomenon by assuming that the anomaly was not spotted because of the particular type of pragmatic processing hearers are involved in. They suggest that the noun *Moses* is shallow-processed in such a way that the accessed

---

5.   That is precisely one of the basic claims of RT: the representation a hearer derives is taken to sufficiently *resemble* the representation the speaker meant to communicate, in terms of its contextual implications (see Wilson 2000).

6.   Contextual effects result from the interaction between new and old information, and are of three types: adding new information to the cognitive environment, strengthening previously held assumptions, and eliminating previously held assumptions (cf. Sperber & Wilson 1995:108–117)

representation is an *ad hoc concept*, *MOSES, which is interpreted as a mere character that fits the biblical context. *MOSES constitutes an underdetermined concept which can refer to a greater set of individuals which include Moses *and* Noah. Crucially, the hearer does not access the concept MOSES: the specific individual who led the people of Israel away from Egypt. In other words, referent saturation turns out to be inadequate, but still ensures a global fit to context of the item in question,[7] and this is precisely why the anomaly is missed by over 2/3 of tested subjects.

We would like to take the analysis of this example a step further in order to understand *why* it is that the concept MOSES is shallow-processed. In the experimental setting (as well as in anecdotal instances of this phenomenon), as noted by Erickson and Mattson (1981: 543) the hearer is *always* requested to answer a question; in other words, it is requested from him that he performs a very specific task — i.e. providing a relevant answer, as opposed, to, say, critically analyse the lexical items contained in the question — which could explain in part why the lexical inconsistency is not spotted, let alone even accessed as anomalous. Furthermore, it could plausibly be assumed that subjects take the nature of the question to be some sort of challenge, where the subject's knowledge is being tested on a particular issue. We would like to argue that the question-answering condition constrains the subject's response, and that in this case the main goal of the exchange goes beyond sense-making: relevance is achieved here not by merely reaching a satisfying interpretation, but by complying with the task the speaker requests from the subject, i.e. by providing a suitable answer. Quoting Sperber and Wilson's definition of cognitive effects, whereby 'a positive cognitive effect is a cognitive effect that contributes positively to the fulfilment of cognitive functions or goals' (1995: 265), we argue that answering the question indeed constitutes, in this case, an overarching cognitive goal, so much so that it can override the goal of making proper sense of the utterance.

Building on the discussion of this example, we propose that manipulative communication is designed to get a hearer to unwittingly shallow-process not just particular concepts, but entire contexts of interpretation. We thus specify here that the central claim of our model is not to say that cognitive processing differs from manipulative to non-manipulative instances of communication; in fact, since the

---

7.   Even if we have no evidence that this plays a role in the explanation, it is worth noting that there is a close semantic vicinity between Moses and Noah: both are old men, both have something to do with water (Moses split the Red Sea in two), and both were depicted as men on whose shoulders great responsibilities rested (saving people vs. saving all animal species). This in turn could be part of the explanation of why Erickson and Mattson (1981) found that replacing Moses with Nixon, or other biblical figures such as Abraham or Joshua, increased the rate of anomaly detection.

success of manipulation lies in its covertness (see Maillat & Oswald 2009), we are led to postulate that manipulated hearers follow the standard processing pattern. And indeed, it is crucial to manipulation that successfully manipulated subjects consider that the utterance they processed *was* relevant in the context they interpreted it against. In short, we claim that manipulation works because our cognitive system provides a possibility of being deceived, by virtue of its fallible nature; a property Sperber et al. (1995) refer to as the human tendency to 'cognitive optimism'.[8]

The rationale of our proposal, therefore, lies in the construal of manipulation as a natural consequence of the risk we have to take when processing information. This means that the model we propose does not differ significantly from a cognitive pragmatic model of communication; accordingly, we aim at specifying the conditions under which cognitive processing may be misled to fulfil the speaker's manipulative intention.

4.   Manipulation as contextual constraint

We posit that manipulation is a twofold process: on the one hand, it induces the hearer into processing the information in a very constrained context of interpretation, and, on the other hand, it simultaneously makes sure that the hearer is prevented from expanding the latter, so that further assumptions (e.g. about the utterance's tentative incompatibility with previously held beliefs, or about the speaker's motivations) are not accessed at all. We thus characterise manipulation in the following way:

> Manipulative communication is a twofold process by which a context-selection constraint is combined with the target utterance in order to (i) force its interpretation within a limited context, and (ii) effectively block access to any alternative contextual assumptions.

4.1   Constraining the context: Salience of contextual assumptions

The first aspect of manipulative communication thus lies in getting the hearer to shallow-process the target utterance in a highly constrained context. That means that in order to be successful, the manipulator must make sure that the hearer

---

8.   '[…] [P]eople are nearly-incorrigible 'cognitive optimists'. They take for granted that their spontaneous cognitive processes are highly reliable, and that the output of these processes does not need re-checking' (1995: 90).

selects only certain assumptions so as to form a subset of his cognitive environment in which the target utterance is relevant, the idea being that the selection of such a subset can be brought about, or made readily accessible, through different constraining mechanisms.

This calls first for a brief description of standard contextual selection, and we will rely on Sperber and Wilson's model in order to describe this procedure.

A first feature of contextual selection is that it is a dynamic process whose output is to be taken as a *variable*; the context of interpretation is not given, but constructed as the linguistic stimuli are processed by the hearer. This means that the set of assumptions that will be selected for this purpose is the result of a cognitive calculation of some sort. RT assumes that this calculation operates by allocating cognitive resources to the mobilisation of relevant information, following the economical constraint evoked above: the less effort it takes to derive a representation, the more it will be relevant; the more cognitive effects a representation yields, the more relevant it will be; and accordingly, the first representation to yield a satisfactory balance between effort and effect will be selected.

A second feature of contextual selection is that it yields representations which adequately comply with this effort/effect trade-off. In other words, contextual selection privileges *salient* representations: that means that *accessible* assumptions as well as *strong* assumptions (in terms of the degree of epistemic validity one is likely to ascribe them) will be privileged, since accessibility and strength contribute respectively to highly probable selection and a high probability of validity. Furthermore, it is reasonable for the hearer to assume that the assumptions he considers salient correspond to those the speaker had in mind at the time of her utterance by virtue of the default presumption of optimal relevance which according to RT governs standard, cooperative, linguistic exchanges.[9] Under this view, both the accessibility and the strength of assumptions (which determine their salience) may contribute to granting them relevance, thereby increasing their chance of being selected in the interpretative process. Moreover, the inclusion of contextual assumptions for the purpose of comprehension is not an open-ended procedure; it is assumed to stop once a level of optimal relevance has been reached.

---

9. The presumption of optimal relevance is formulated as follows (cf. Sperber & Wilson 1995:267):

   (a)   The set of assumptions *I* which the communicator intends to make manifest to the addressee is relevant enough to make it worth the addressee's while to process the ostensive stimulus.
   (b)   The ostensive stimulus is the most relevant one the communicator could have used to communicate *I*.

On a more technical note, a context is a set of mental representations which are used as premises in the comprehension procedure. These representations may be *about* very different things: social relationships, events, concepts, linguistic items, mental states, general knowledge, perceptual evidence, etc. What determines their inclusion in the context of interpretation is their relevance towards the utterance at stake, i.e. the amount of cognitive efforts their derivation requires and the effects they are expected to yield with respect to the efficiency of the comprehension procedure. Accordingly, salient assumptions, that is to say both easily accessible and epistemically strong ones, will stand higher chances to be included during the context selection process. Context selection is, thus, regarded as taking place in a structured matrix of contextual assumptions in which the organisational principle is that of salience as shown in Figure 1.
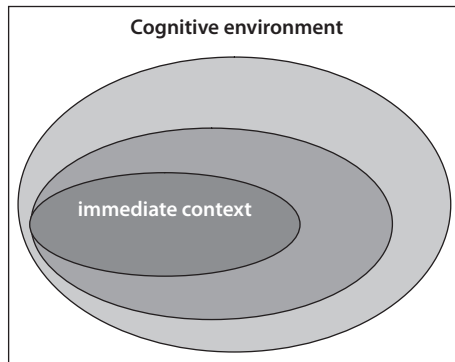

Figure 1. Context selection functions incrementally

The view represented in Figure 1 is reflected in Sperber and Wilson's account, in which contextual selection is assumed to be an incremental procedure, that is, that it starts off from a minimal, initial, context of cognitively 'cheap' assumptions (such as for instance assumptions about the immediate observable environment, or assumptions which result from immediate previous processing), and gradually adds more 'costly' assumptions to the context set, in order to expand the latter until it reaches a level of optimal relevance. In other words, context selection is a partially ordered process based on an inclusion relation. Interestingly, Sperber and Wilson consider that, from a psychological perspective, 'order of inclusion corresponds to order of accessibility' (1995: 142), meaning that the more accessible an assumption will be, the more likely it will be selected in the context. In Figure 1, it corresponds to the most salient assumptions being closer to the innermost subset (i.e. the 'immediate context').

Let us now return to the question of manipulative discourse. The general procedure carries a straightforward implication for the success of a manipulative attempt, such as defined in (1) above: the more a manipulator will be successful in making assumptions salient, the more these will be likely to be part of the final context of interpretation, and the more she will be likely to be successful in fulfilling her manipulative intention.

Specifically, trying to sum up the main claim of this chapter, we argue that manipulative communication is a twofold process by which a manipulative operation constrains the context selection process of a target utterance, which will eventually lead to the shallow processing of the latter within the restricted context. Manipulation is therefore an attempt at controlling the context selection process of an utterance U by making a set of assumptions C so salient so as to make them inescapable from a cognitive point of view (see Maillat 2006). In accordance with the relevance-theoretic framework, it is predicted that the context-selection constraint can essentially take two forms. The manipulator can try to make some contextual assumptions more accessible, which constitutes an effort-oriented strategy; or she can try to make these contextual assumptions stronger, which would be an effect-oriented strategy. Crucially, the inescapability of C is intended to block access to a context C' in which U would be incompatible with other contextual assumptions.

If we now turn to research in (cognitive) psychology, but also, interestingly enough, to the paradigm of Argumentation Theory, there is a vast literature that contributes to explaining how assumptions can be made highly salient so as to become 'cognitively inescapable'. We already mentioned the task-constraining feature of the *Moses illusion*, which induces the addressee into achieving relevance out of the production of a specific answer. Some illusions, such as the *Validity effect* (see Hackett Renner 2004), also known as the effect of repetition, illustrate how accessibility may lead to validity (and strength). The idea, which is often exploited in advertising techniques, is to repeat over and over the same message. It has been experimentally shown that the degree of validity one ascribes to a message increases through its repetition, regardless of the initial degree of validity one ascribed to it (that is, whether the statement was initially believed to be true or false is irrelevant). We take this experimentally documented psychological feature to be one technique of contextual constraint that manipulators may exploit; repetition increases validity, which has been shown to correspond to an increase in cognitive strength, making the assumption an ideal candidate for later inclusion in the context set of a target utterance.[10]

---

10.   See below for an illustration of contextual constraining of accessibility with real-life examples.

An illustration of ways to increase the strength of assumptions is to be found in the phenomena described in the paradigm of Argumentation Theory, namely fallacies, that is, argumentative schemes which happen to violate specific norms of soundness.

5.   A Cognitive insight into fallacies

Let us take the case of the *ad verecundiam* fallacy (see e.g. van Eemeren et al. 2002: 131) which relies on resorting to the voice of an expert in order to present an argument as unarguable. This type of fallacy is traditionally taken to enter the general inventory of rhetoric devices used for manipulative purposes. We are now in a position to provide a pragmatic analysis of its cognitive underpinnings. We shall illustrate this by means of a recent case of manipulation which received a great deal of media exposure. In 2008, some journalists brought up evidence that the military analysts working as independent *experts* for various major American TV channels were being briefed by the Pentagon about what and how they should evaluate the tactics, events, reports provided by the US Military about its ongoing military operations in various theatres of conflict, in particular Iraq and Afghanistan. In other words, the Pentagon was trying to control the voice of the expert.[11]

If we go back to our model, this type of manipulation lends itself to a straightforward pragmatic analysis. As we pointed out earlier, a first manipulative discursive device puts a constraint on the context selection process of a target utterance whose interpretation is thus restricted to a limited, crucially incomplete, set of contextual assumptions. In our example, the constraining mechanism involves a complex manoeuvre whereby the manipulator controls the assumptions provided by the experts. The interest of such a strategy lies in the added salience that expert's assumptions have in the cognitive environment of the hearer due to their greater strength, i.e. their heightened validity. By controlling the experts, the manipulator is able to guarantee that a given, constrained subset of highly salient assumptions will be available to her designated audience when they have to process news regarding e.g. American casualties in Iraq or the treatment of prisoners held in Guantánamo.

As we hinted above, from this particular case study we can generalise our analysis to *ad verecundiam* fallacies. Such fallacies are efficient and part of traditional manipulative techniques because they cover a range of strategies which

---

11.   See the *New York Times*' coverage of the scandal (http://www.nytimes.com/2008/04/20/washington/20generals.html).

take advantage of the cognitive loopholes in the inferential mechanism governing manipulation, as was explained in the previous section. Specifically, the *ad verecundiam* fallacies exploit the fact that not all assumptions are born equal. Instead, some assumptions carry more contextual weight. They are more salient and therefore more accessible during context selection. This particular cognitive bias can be exploited in a variety of ways: by controlling the expert's voice as was the case in our example; by assigning expertise to some specific speaker; or by confusing fame with authority.[12]

A second example will allow us to get a better grasp of the kind of insight provided by our model into manipulative discourse. Another well-known strategy used by manipulators is based on the *ad populum* fallacy, whereby an argument is presented as stronger since 'everybody says so' (see e.g. van Eemeren et al. 2002:131). In this second instance we will not resort to a real-life example but we will simply refer back to the empirical support we highlighted earlier in connection with the so-called *validity effect*. We showed that cognitive psychologists (see Hacket Renner 2004) had experimentally established that there exists a direct positive correlation between the degree to which a given assertion is repeated and the degree of validity assigned to it by subjects. Essentially, this means that the more often a hearer comes across a given statement the more valid this statement will appear to be for him. This provides empirical support for the cognitive analysis of manipulation we advocate in explaining well-known advertising techniques, as well as methods applied in ill-famed propaganda strategies. It is worth pointing out that in this instance, manipulation takes advantage of a combination of cognitive biases. First, there is the psychological bias which assigns higher validity to contextual assumptions that have a higher frequency. Second, there is the human tendency to cognitive optimism identified by Sperber, Cara and Girotto (1995) which leads hearers to take cognitive shortcuts when they favour more salient contextual assumptions in their quest for optimal relevance.

While this last point is worth mentioning in its own right, it acquires even more weight when we consider it in connection with the *ad populum* fallacy. In fact, if we are right in the analysis we make of the manipulative power of the *validity effect* — and there is strong empirical support that this is indeed the right approach — this latter type of fallacious argument follows quite naturally. The well-attested

---

12.   The latter is not — strictly speaking — an instance of *ad verecundiam* fallacy. It would include the nowadays common political ploy which consists in staging a meeting with some rock star who will vouch for her unconditional support of the presidential candidate. However, it is interesting to see that our model provides a new cognitive basis for a taxonomy of manipulative techniques.

efficiency of the *ad populum* fallacy is a direct consequence of our cognitively experiencing the validity effect. Effectively, if 'everybody says so', a given assumption is repeated and will inevitably be more valid, therefore stronger, i.e. more salient; and finally following our account, inescapable when constructing the set of contextual assumptions of the target utterance. In other words, while the *ad populum* argument might be fallacious on an argumentative level, it is a cognitively grounded fact that assumptions which are heard repeatedly will have a higher validity, and will therefore be more likely to be selected.

This second example aptly shows how the proposed model sheds new light on argumentative phenomena. The cognitive basis of the proposal and its connection with the empirical research carried out in the adjacent field of cognitive psychology provides us with new tools to explain some well-known manipulative ploys while allowing us to explore and redefine the domain. Strikingly, the suggested framework is also capable of capturing manipulative techniques which, while they clearly belong to the domain of manipulation, do not lend themselves easily to any straightforward analysis within any of the more traditional approaches (see the discussion in Saussure & Schulz 2005).

6.   Adding some spin to the model

On 11 September 2001, shortly after two planes had been flown into the World Trade Centre in New York, Jo Moore, the then special adviser to the British Secretary of State, Stephen Byers, circulated an email in her service pointing out that that day 'was a good day to bury bad news'.[13] She was hinting at a difficult announcement that needed to be made about some expenses incurred by councillors. In doing so, she was probably only doing her job as a spin doctor helping the government to handle the media. Interestingly for our purpose, this common and transparent form of manipulation provides an excellent means to illustrate the explanatory power of the proposed model.

In our terms, the well-known strategy used by Moore in this instance is a manipulative attempt at constraining contextual selection to a context of *maximal irrelevance*, as it were, for the target utterance (in this case the statement regarding the councillors' expenses). The strength of our analysis lies in its ability to straightforwardly account for this phenomenon in spite of its problematic nature. The problem encountered by other analytic frameworks with an

---

13.   The English newspaper *The Guardian* extensively covered the affair (see http://www.guardian.co.uk/media/2001/oct/09/terrorismandthemedia.Whitehall).

example like this lies in the fact that the target utterance is not misleading in any obvious way — a true statement is made about some disturbing fact in a transparent way. The essence of a manipulative strategy, however, does not reside in the target utterance, but as we pointed out earlier in the first operation which constrains context selection. In this instance, it corresponds to the decision to select a cognitive environment which includes the 9/11 events. Again, the literature in cognitive psychology gives us some very clear clues as to why such a context might help the spin doctor in achieving her manipulative goal. According to Pickrell et al. (2004: 352–353), who use 9/11 as a case in point, some contexts constitute 'highly salient memories' which have a 'highly emotional, meaningful, and subjectively permanent nature'. They have been called 'flashbulb' contexts. In other words, some contexts are cognitively overwhelming and, therefore, inescapable. This property renders them extremely useful when it comes to constraining context selection during the first manipulative operation because they can cognitively erase the target utterance as it were. In this respect, they could be described as *cognitive black holes* from a relevance-theoretic viewpoint.

Specifically, in such a case, the overwhelming salience of the subset of contextual assumptions associated with 9/11 is such that the processing of any unrelated utterance will only achieve very low relevance in the hearer's cognitive environment. As a consequence, the assumptions retrieved during the processing of that utterance will promptly vanish, which is precisely the manipulative effect sought by the spin doctor. As Moore herself puts it, 9/11 was, at the time, cognitive quicksand in which bad news were buried.

7. Conclusion

What we see with these illustrations is that context selection can be constrained in a variety of ways. Crucially, the very variety of manipulative strategies which we have only briefly touched upon in this chapter is laid out in and predicted by the theoretical model itself. The resulting typology of manipulative techniques can be demonstrated to range over the various tactics which can affect contextual salience, either by making a contextual subset more accessible (effort-based) or by making it stronger (effect-based). In many cases, these types correspond to existing argumentative fallacies or infamous manipulative ploys, while in some other cases — like that of cognitive black holes — the cognitively-based pragmatic model described above opens new perspectives in our understanding of these forms of deceitful communication.

In this chapter, we have shown that very different manipulative strategies can be captured through a unified, pragmatic model which, most interestingly, finds independent support in the work carried out in another domain. One of the very promising aspects of this research lies in the inherent testability of its main hypothesis within an experimental setting, as well as in the wealth of existing empirical designs which have paved the way for our investigation of cognitive optimism; this very human flaw which lies at the heart of manipulative discourse.

References

Allott, N. (2008). *Pragmatics and Rationality*. Ph.D. Thesis, available at http://www.phon.ucl.ac.uk/home/nick/content/n_allott_phd_thesis_09_2007.pdf.

Allott, N. & P. Rubio Fernández (2002). This paper fills a much-needed gap. In P. Afuta, A. El Ghali & F. Toussenel (eds.), *Actes de l'Atelier des Doctorants en Linguistique.* Université Paris *7*. pp. 97–102.

Erickson, T.A. & M.E. Mattson (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior* 20: 540–552.

Fowler, R., R. Hodge, G. Kress & T. Trew (1979). *Language and Control.* London: Routledge and Kegan Paul.

Grice, H.P. (1989). *Studies in the Way of Words*. Harvard: Harvard University Press.

Hackett Renner, C. (2004). Validity Effect. In R.F. Pohl (ed.), *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgment and Memory*. New York: Psychology Press. pp. 201–213.

Maillat, D. (2006). Is manipulation relevant? Paper presented at University della Svizzera Italiana, Lugano.

Maillat, D. & S. Oswald (2009). Defining manipulative discourse: The pragmatics of cognitive Illusions. *International Review of Pragmatics* 1 (2): 348–370.

Park, H. & L. Reder (2004). Moses illusion. In R.F. Pohl (ed.), *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgment and Memory*. New York: Psychology Press. pp. 275–291.

Pickrell, J., D. Bernstein & E. Loftus (2004). Misinformation effect. In R. Pohl (ed.), *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*. Hove & New York: Psychology Press. pp. 345–361.

Pohl, R. (ed.) (2004). *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgment and Memory*. New York: Psychology Press.

Rigotti, E. (2005). Towards a typology of manipulative processes. In L. de Saussure & P. Schulz (eds.), *Manipulation and Ideologies in the Twentieth Century: Discourse, Language, Mind.* Amsterdam: John Benjamins. pp. 61–83.

Saussure, L. de & P. Schulz (eds.) (2005). *Manipulation and Ideologies in the Twentieth Century: Discourse, Language, Mind*. Amsterdam: John Benjamins.

Sperber, D., & D. Wilson (1995). *Relevance: Communication and Cognition* (2nd ed.). Oxford: Blackwell.

Sperber, D., F. Cara, & V. Girotto (1995). Relevance Theory explains the Selection Task. *Cognition* 57: 31–95.

Tversky, A. & D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science* 185(4157): 1124–1131.

Van Dijk, T. (2006). Discourse and manipulation. *Discourse & Society* 17 (3): 359–383.

Van Eemeren, F., Grootendorst, R. & F. Snoeck Henkemans (2002). *Argumentation: Analysis, Evaluation, Presentation.* London: Routledge.

Wilson, D.(2000). Metarepresentation in linguistic communication. In D. Sperber (ed.), *Metarepresentations: A Multidisciplinary Perspective.* Oxford: Oxford University Press. pp. 411–448.