Dan Sperber and Ira Noveck. 2004. Introduction. in *Experimental Pragmatics*, Noveck, I. & Sperber, D. (eds). Palgrave.

This volume lays down the bases for a new field, Experimental Pragmatics, that draws on pragmatics, psycholinguistics and also on the psychology of reasoning.

Dan Sperber and Ira Noveck.
Introduction. to ***Experimental Pragmatics***

How does our knowledge of language on the one hand, and of the context on the other permit us to understand what we are told, resolve ambiguities, grasp both explicit and implicit content, recognize the force of a speech act, appreciate metaphor and irony? These issues have been studied in two disciplines: pragmatics and psycholinguistics, with limited interactions between the two. Pragmatics is rooted in the philosophy of language and in linguistics and has spawned competing theories using as evidence a mixture of intuitions about interpretation and observations of behavior. Psycholinguistics has developed sophisticated experimental methods in the study of verbal communication, but has not used them to test systematic pragmatic theories. This volume lays down the bases for a new field, Experimental Pragmatics, that draws on pragmatics, psycholinguistics and also on the psychology of reasoning. Chapters in this volume either review pioneering work or present novel ways of articulating theories and experimental methods in the area. In this introduction we outline some core pragmatic issues and approaches and relate them to experimental work in psycholinguistics and in the psychology of reasoning. We then briefly present one by one the chapters of this collection.

1. Some core pragmatic issues and approaches

In a very broad sense, pragmatics is the study of language use. It encompasses loosely related research programmes ranging from formal studies of deictic expressions to sociological studies of ethnic verbal stereotypes. In a more focused sense, pragmatics is the study of how linguistic properties and contextual factors interact in the interpretation of utterances. We will be using “pragmatics” only in this narrower sense. Here we briefly highlight a range of closely related, fairly central pragmatic issues and approaches that have been of interest to linguists and philosophers of language in the past thirty years or so, and that, in our opinion, may both benefit from, and contribute to, work in experimental psychology.

A sentence of a language can be considered as an abstract object with phonological, syntactic and semantic properties assigned by the grammar of the language (the grammar itself being generally seen as a mental system). The study of these grammatical properties is at the core of linguistics. An utterance, by contrast, is a concrete object with a definite location in time and space. An utterance is a realisation of a sentence (a realisation that can be defective in various respects, for instance by being mispronounced). An utterance inherits the linguistic properties of the sentence it realises and has further properties linked to its being uttered in a given situation by a speaker addressing an audience. In verbal communication, both linguistic and non-linguistic properties of utterances are involved. But what role exactly do these properties play and how do they interact? These are questions that pragmatic theories attempt to answer.

The pragmatic approaches we are concerned with here all accept as foundational two ideas that have been defended by the philosopher Paul Grice (Grice, 1989) . The first idea is that, in verbal communication, the interlocutors share at least one goal: having the hearer recognize the speaker's meaning. The linguistic decoding of the sentence uttered provides the hearer with the sentence meaning, but this decoding is only a subpart of the process involved in arriving at a recognition of the speaker's meaning. This recognition does not involve any distinct awareness of the sentence meaning, that is, of the semantic properties assigned to the sentence by the grammar. Only linguists and philosophers of language have a clear and distinct notion of, and an interest in, sentence meaning proper. Unlike sentence meaning, which is an abstraction, a speaker's meaning is a mental state. More specifically, for a speaker to mean that P is for her to have the intention that the hearer should realise that, in producing her utterance, she intended him to think that P. A speaker's meaning is an overt intention that is fulfilled by being recognised by the intended audience. Consider for instance Mary's contribution to the following exchange:

(1) Peter: Do you like Fellini's films?

Mary: Some of them.

In replying "some of them," Mary intends Peter to realize that she intends him to think that she likes some of Fellini's films, but not all. The proposition Mary likes some of Fellini's films but not all is Mary's meaning. It is not the linguistic meaning of the sentence fragment "some of them," which can be used in other situations to convey totally different contents. Mary's meaning goes well beyond the meaning of the linguistic expression she uttered.

Verbal comprehension is often seen in psycholinguistics as the study of linguistic decoding processes, drawing on grammar (with the possibility that grammar may extend above the level of the sentence to that of discourse) and using contextual factors in a limited way, to disambiguate ambiguous expressions and fix reference. The idea that successful communication consists in the recognition by the audience of the speaker's meaning suggests a different approach. Verbal comprehension should be seen as a special form of attribution of a mental state to the speaker. This attribution is dependent on linguistic decoding, but is essentially an inferential process using as input the result of this decoding and contextual information.

The second foundational idea defended by Paul Grice is that, in inferring the speaker's meaning on the basis of the decoding of her utterance and of contextual information, the hearer is guided by the expectation that the utterance should meet some specific standards. The standards Grice envisaged were based on the idea that a conversation is a cooperative activity. Interlocutors are expected to follow what he called a "co-operative principle" requiring that they "make [their] conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which [they] are engaged." This is achieved by obeying a number of "maxims of conversation" which Grice expressed as follows:

Maxims of Quantity:

1. Make your contribution as informative as is required (for the current purpose of the exchange)

2. Do not make your contribution more informative than is required

Maxims of Quality:

Supermaxim: Try to make your contribution one that is true.

1. Do not say what you believe to be false.

2. Do not say that for which you lack adequate evidence.

Maxim of Relation:

Be relevant.

Maxims of Manner:

Supermaxim: Be perspicuous

1. Avoid obscurity of expression.

2. Avoid ambiguity.

3. Be brief (avoid unnecessary prolixity).

4. Be orderly.

In interpreting an utterance, the best hypothesis for the hearer to choose is the one that is the most consistent with the assumption that the speaker has indeed followed these maxims. For instance, in interpreting Mary's reply "some of them" in the above dialogue, Peter is entitled to draw several inferences. He is entitled, in the first place, to treat this sentence fragment as elliptical for "I like some of Fellini's films" since this is the interpretation most consistent with the assumption that Mary was following the maxims, and in particular the maxims "be relevant" and "be brief." Peter is also entitled to understand Mary to mean that she does not like all of Fellini's films. If she did like all of them, she would be violating the maxim "make your contribution as informative as is required" in talking only of "some of them."

Current pragmatic theories draw on Grice's idea that the existence of set expectations is what allows hearers to infer the speaker's meaning on the basis of the utterance and the context. These theories differ in their account of the precise expectations that drive the comprehension process. Neo-Griceans (Atlas, forthcoming; Gazdar, 1979; Horn, 1973 , 1984, 1989, 1992; Levinson, 1983, 2000) stay relatively close to Grice's formulation. Levinson (2000), for instance, defines three basic principles linked to three of Grice's maxims (here in abridged form):

Q-Principle

Speaker’s maxim: Do not provide a statement that is informationally weaker than your knowledge of the world allows.

Recipient corollary: Take it that the speaker made the strongest statement consistent with what he knows.

I-Principle

Speaker’s maxim: Produce the minimal linguistic information sufficient to achieve your communicational ends.

Recipient corollary: Amplify the informational content of the speaker’s utterance, by finding the most specific interpretation, up to what you judge to be the speaker’s […] point.

M-Principle

Speaker’s maxim: Indicate an abnormal, nonstereotypical situation by using marked expressions that contrast with those you would use to describe the corresponding normal, stereotypical situations.

Recipient corollary: What is said in an abnormal way indicates an abnormal situation.

These principles provide heuristics for interpreting utterances. For instance, when Mary answers elliptically “some of them”, she can be seen by Peter as producing the minimal linguistic information sufficient to achieve her communicational ends (following the I Principle), and this, together with the assumption that Mary obeyed the Gricean Maxim of relevance, justifies his amplifying the content of her utterance up to what he judges to be her point (see Levinson 2000, 183-4). Moreover, the Q Principle justifies Peter in taking it that Mary made the strongest statement consistent with her knowledge, and that therefore it is not the case that she likes all of Fellini’s films.

Relevance theory (Bezuidenhout, 1997; Blakemore, 1987, 2002; Blass, 1990; Carston, 2002; Carston & Uchida, 1997; Gutt, 1991 Ifantidou, 2001; Matsui, 2000; Moeschler, 1989; Noh, 2000; Papafragou, 2000; Pilkington 2000; Reboul, 1992; Rouchota & Jucker, 1998; Sperber & Wilson, 1986/1995; Yus, 1997), though still based on Grice’s two foundational ideas, departs substantially from his account of the expectations that guide the comprehension process. For Griceans and neo-Griceans, these expectations derive from principles and maxims, i.e., rules of behaviour that speakers are expected to obey but may, on occasion, violate. Such violations may be unavoidable because of a clash of maxims or of principles, or they may be committed on purpose in order to indicate to the hearer some implicit meaning. Indeed, in the Gricean scheme, the implicit content of an utterance is typically inferred by the hearer in his effort to find an interpretation which preserves the assumption that the speaker is obeying, if not all the maxims, at least the co-operative principle. For Relevance Theory, the very act of communicating raises in the intended audience precise and predictable expectations of relevance, which are enough on their own to guide the hearer towards the speaker’s meaning. Speakers may fail to be relevant, but they may not, if they are communicating at all (rather than, say, rehearsing a speech), produce utterances that do not convey a presumption of their own relevance.

Whereas Grice invokes relevance (in his "maxim of relation") without defining it at all, Relevance Theory starts from a detailed account of relevance and its role in cognition. Relevance is defined as a property of inputs to cognitive processes. These inputs include external stimuli, which can be perceived and attended to, and mental representations, which can be stored, recalled, or used as premisses in inference. An input is relevant to an individual when it connects with background knowledge to yield new cognitive effects, for instance by answering a question, confirming a hypothesis, or correcting a mistake. Slightly more technically, cognitive effects are changes in the individual's set of assumptions resulting from the processing of an input in a context of previously held assumptions. This processing may result in three types of cognitive effects: the derivation of new assumptions, the modification of the degree of strength of previously held assumptions, or the deletion of previously held assumptions. Relevance, i.e. the possibility of achieving such a cognitive effect, is what makes an input worth processing. Everything else being equal, inputs which yield greater cognitive effects are more relevant and more worth processing. For instance, being told by the doctor "you have the flu" is likely to carry more cognitive effects and therefore be more relevant than being told "you are ill." In processing an input, mental effort is expended. Everything else being equal, relevant inputs involving a smaller processing effort are more relevant and more worth processing. For instance, being told "you have the flu" is likely to be more relevant than being told "you have a disease spelled with the sixth, the twelth, and the twenty-first letter of the alphabet" because the first of these two statements would yield the same cognitive effects as th second for much less processing effort. Relevance is a thus matter of degree and varies with two factors, positively with cognitive effect, and inversely with processing effort.

Relevance Theory develops two general claims or "principles" about the role of relevance in cognition and in communication:

Cognitive Principle of Relevance: Human cognition tends to be geared to the maximisation of relevance

Communicative Principle of Relevance: Every act of communication conveys a presumption of its own optimal relevance

As we have already mentioned, these two principles of relevance are descriptive and not normative (unlike the principles and maxims of Gricean and neo-Gricean pragmatists). The first, Cognitive Principle of Relevance yields a variety of predictions regarding human cognitive processes. It predicts that our perceptual mechanisms tend spontaneously to pick out potentially relevant stimuli, our retrieval mechanisms tend spontaneously to activate potentially relevant assumptions, and our inferential mechanisms tend spontaneously to process them in the most productive way. This principle, moreover, has essential implications for human communication processes. In order to communicate, the communicator needs her audience's attention. If, as claimed by the Cognitive Principle of Relevance, attention tends automatically to go to what is most relevant at the time, then the success of communication depends on the audience taking the utterance to be relevant enough to be to be worthy of attention. Wanting her communication to succeed, the communicator, by the very act of communicating, indicates that she wants her utterance to be seen as relevant by the audience, and this is what the Communicative Principle of Relevance states.

According to Relevance Theory, the presumption of optimal relevance conveyed by every utterance is precise enough to ground a specific comprehension heuristic:

Presumption of optimal relevance:

(a) The utterance is relevant enough to be worth processing;

(b) It is the most relevant one compatible with communicator's abilities and preferences.

Relevance-guided comprehension heuristic:

(a) Follow a path of least effort in constructing an interpretation of the utterance (and in particular in resolving ambiguities and referential indeterminacies, in going beyond linguistic meaning, in computing implicatures, etc.).

(b) Stop when your expectations of relevance are satisfied.

For instance, when Mary, in response to Peter's question "Do you like Fellini's films?" utters "some of them," she can be confident that, following a path of least effort, Peter will understand "them" to refer to Fellini's films (since this is the plural referent most prominent in his mind) and the whole utterance to be elliptical for "I like some of them" (since this is the resolution of the ellipsis closest to his expectations). The fact that there are films by Fellini that Mary likes is relevant enough to be worth Peter's attention (as he indicated it would be by asking the question). However, this does not yet fully satisfy Peter's expectations of relevance: Mary was presumably able, and not reluctant, to tell him whether she liked all of Fellini's films, and that too would be of relevance to Peter. Given that she did not say that she likes them all, Peter is entitled to understand her as meaning that she likes only some of them. Having so constructed the interpretation of Mary's utterance, Peter's expectations of relevance are now satisfied, and he does not develop the interpretation any further.

Grice's original theory, the Neo-Gricean theory and Relevance Theory are not the only theoretical approaches to pragmatics (even in the restricted sense of "pragmatics" we adopt here). Important contributors to pragmatic theorizing with original points of view include Anscombre & Ducrot (1995); Bach (1987, 1994); Bach & Harnish (1979); Blutner & Zeevat (2003); Dascal (1981); Ducrot, (1984); Fauconnier (1975, 1985); Harnish (1976, 1994); Kasher (1976, 1984, 1998); Katz (1977); Lewis (1979); Neale (1990, 1992, forthcoming); Recanati (1979, 1988, 1993, 2000); Searle (1969, 1979); Stalnaker (1999); Sweetser (1990); Travis (1975); Van der Auwera, J. (1981, 1985, 1997); Vanderveken (1990-91); see also Davis (1991), Moeschler & Reboul (1994). However the three approaches we have briefly outlined here are arguably the dominant ones, and the most relevant ones to the experimental research reported in this book.

2. What can pragmatic theories and experimental psycholinguistics offer each other?

Griceans, Neo-Griceans, Relevance Theorists, and other pragmatists, all have ways to account for examples such as (1) above, and for pragmatic intuitions generally. It is hard to find in pragmatics crucial evidence that would clearly confirm one theory and disconfirm another. To experimental psychologists, it might be obvious that one should use experimental evidence in order to evaluate and compare pragmatic claims. Pragmatics, however, has been developed by philosophers of language and linguists who often have little familiarity with, or even interest in, experimental psychology. The only source of evidence most of them have ever used has been their own intuitions about how an invented utterance would be interpreted in a hypothetical situation. Provided that intuitions are systematic enough across subjects, there is nothing intrinsically wrong in using them as evidence, as the achievements of modern linguistics (which relies heavily on such intuitions) amply demonstrate. More sociologically oriented pragmatists have insisted on the use of evidence from recordings of genuine verbal exchanges, or of genuine written texts, together with data about the speakers or authors and the situation. Even though the interpretation of these naturally occurring utterances is normally left to the pragmatist's intuitive interpretive abilities, their use has been of great value in investigating a variety of pragmatic issues.

Pragmatic research is not to be censured, let alone discarded, on the ground that it is mostly based on intuition and observational data and has hardly been pursued at all as an experimental discipline. However, this has meant that preference for one theory over another is justified not in terms of crucial empirical tests but mostly on grounds of consistency, simplicity, explicitness, comprehensiveness, explanatory force, and integration with neighbouring fields. For example, it has been argued that Grice's own formulation of his principle and maxims is too vague, and not explanatory enough: Gricean explanations are more like ex post facto rationalisations. Neo-Griceans are developing an approach to pragmatics in close continuity with linguistic semantics, and view this as an advantage. Relevance theorists feel that their approach is more explanatory, more parsimonious, and better integrated into the cognitive sciences. These considerations, however relevant to evaluating theories, can themselves be diversely evaluated.

Turning from pragmatics to experimental psycholinguistics—an older and more developed science—, we find a rich and extensive domain of research dealing with diverse themes ranging from the child's first language acquisition to the mechanics of speech production. Among these themes is that of comprehension, which includes a variety of sub-themes from the perception and decoding of the acoustic (or visual) signal to the interpretation of discourse. In principle, the range of phenomena that pragmatics investigates is part of the much wider domain of psycholinguistics. However, with its own rich history, traditions, and focus on experimental research, psycholinguistics has generally paid very little attention to the discipline of pragmatics, even when the phenomena studied have been standard pragmatic phenomena. Rather, it has developed its own theoretical approaches to pragmatic themes, in particular under the label of 'discourse processes.' To what extent, and on what specific points research on discourse processes might converge or conflict with specific pragmatic claims remains largely to be seen (for a comparison between the psycholinguistic notion of discourse coherence and the pragmatic notion of relevance see in particular Blakemore 2001, 2002, Blass 1990, Rouchota 1998, Unger 2000, Wilson 1998, Wilson & Matsui 2000).

It is reasonable to expect that two fields of research dealing in part with the same material at the same level of abstraction would gain by joining forces, or at least by interacting actively. For pragmatics the gain would be twofold. First, experimental evidence can be used, together with intuition and recordings, to confirm or disconfirm hypotheses. The high reliability and strong evidential value of experimental data puts a premium on this sort of data even though it is hard to collect and is generally more artificial than observational data (and therefore raises specific problems of interpretation). The three kinds of evidence—intuitions, observations and experiments—are each in their own way relevant to suggesting and testing pragmatic hypotheses, and they should be used singly or jointly whenever useful. Secondly, aiming at experimental testability puts valuable pressure on theorizing. Too often, armchair theories owe much of their appeal to their vagueness, which allows one to reinterpret them indefinitely so as to fit one's understanding of the data, but which also makes them untestable. Developing an experimental side to pragmatics involves requiring a higher degree of theoretical explicitness. Moreover, experimentally testing theories often leads one to revise and refine them in the light of new and precise evidence, and gives theoretical work an added momentum.

For experimental psycholinguistics, the gain from a greater involvement with pragmatics would be in taking advantage of the competencies, concepts and theories developed in this field, in order to better describe and explain a range of phenomena that are clearly of a psycholinguistic nature, and to develop new experimental paradigms. The experimental approach often results in imbalanced coverage of the domain of study. Topics for which an experimental paradigm has been developed get studied in great detail, whereas other topics of comparable empirical importance may remain largely untouched for lack of an ad hoc experimental tradition. There is, for instance, a wealth of psycholinguistic research on metaphor but very little on implicatures, when, from a pragmatic point of view, the two phenomena are of comparable importance. Typically, pragmatic theories have been more comprehensive and evenly detailed than psycholinguistic ones.

The small amount of existing Experimental Pragmatic work from psycholinguists and pragmatists already shows what this collection is meant to demonstrate, namely that there is much to gain, both for pragmatics and for psycholinguistics, from systematically putting pragmatic hypotheses to the experimental test. Here we give a brief account of two examples: indirect speech acts and bridging.

An early illustration of the relevance of experiments to theoretical issues was provided by experimental work done in the 70s on a topic of hot theoretical debates at the time: indirect speech acts (Searle 1975). When a speaker says "Could you stop fidgeting?" is the speech act a question, or is it a request? The problem with categorizing this as a question is that, in ordinary circumstances, the proper response for the hearer is not to provide a verbal answer such as: "yes, I could" or "no, I couldn't" (as would be appropriate in response to a genuine question), but to actually stop fidgeting. The problem with categorizing it as a request is that the mood of the sentence is interrogative and not imperative. Sentential moods, it is generally assumed, indicate the kind of speech acts an utterance can be used to perform: declaratives serve to make assertions, interrogatives to ask questions, imperatives to make requests, and so forth. Indirect speech acts are called "indirect" precisely because they don't seem to conform to the indication given by their mood: a declarative utterance may indirectly express a request (e.g. "you could stop fidgeting") or a question (e.g. "I would like to know where you have been"), an interrogative utterance may indirectly express a request (as in our example, "could you stop fidgeting?") or an assertion (e.g. "Who could remain indifferent in front of such injustice?"); an interrogative can also serve to ask an indirect wh-question, different from the yes-no question it would express directly (e.g. "could you tell me the time?"), and so on. Indirect speech acts thus seem to threaten a basic assumption of much linguistic thinking. A possible way to go is to treat indirect speech acts as non-literal uses of language, comparable to metaphor and, like metaphor, explainable in pragmatic terms. Another way is to take indirect requests to be conventional or idiomatic. But are these descriptions really adequate? This is where experimental work comes in.

If an indirect speech act is like an idiom with a conventional meaning, then understanding it should not involve more processing than understanding a direct speech act. Reaction time studies, such as those by Clark & Lucy 1975, suggested that, in fact, indirect requests do take longer to comprehend than direct ones and therefore are not conventional (but see Gibbs 1979 for a more complex picture). If indirect requests are like metaphor, then their literal interpretation should not be retained at all. After all, when a sentence is used metaphorically (e.g. "John is a bulldozer"), the literal sense is not at all part of the speaker's meaning. Clark 1979 telephoned store owners with indirect questions such as "Can you tell me what time you close?" and most answered things like "Yes, we close at six." Yes, in such an answer, seems to be an answer to the direct question ("Can you tell me?" "yes, I can tell you") whereas the rest of the sentence ("we close at six") is an answer to the indirect question ("At what time do you close?"), suggesting that both the direct and the indirect questions were considered parts of the speaker's overall meaning (for further evidence and different analyses, see Munro 1979, Gibbs 1981). Not only did these experimental studies provide relevant evidence in the theoretical debate, they also suggested new and more specific hypotheses about indirect speech acts. However, this early dialogue between experimentalists and pragmatists working on indirect speech acts largely ends here. The two groups failed to take as much advantage of each other's work as they could have.

Another example of interactions between psycholinguistics and pragmatics is provided by the case of bridging. A bridging inference, or bridging implicature (Clark, 1973), links a referring expression to an intended referent that is neither present in the environment nor mentioned in the ongoing discourse but that is nevertheless inferentially identifiable. For example, in the two sentences in (2) below:

(2) John walked into a room. The window was open.

The expression the window is a referring expression implicitly linked to the room mentioned in the preceding sentence. In order to establish the link, a bridging implicature such as the room had a window has to be retrieved. Bridging was the basis for one of the first innovative accounts of discourse from Clark and colleagues – the Given-New contract – which has inspired much valuable experimental work in psycholinguistics. This research has contributed to the development of innovative paradigms (e.g. using reading times and semantic probes), for the creation of typologies in texts (Sanders, Spooren, & Noordman, 1992) , and has fed theoretical debates (e.g. between the Constructionist vs. Minimalist accounts of inference generation) in the psychological literature (Graesser, Singer, & Trabasso, 1994; McKoon & Ratcliff, 1992) .

Although Clark explicitly drew inspiration from Grice, and although bridging is obviously an important pragmatic topic, the exchanges between the pragmatic and psycholinguistic communities on the theme of bridging remained limited. A recent exception is provided by the work of Tomoko Matsui (Matsui, 2000) , a pragmatist who has become involved in experimental research. She makes a distinction between cases of bridging proper, like (2), where "contextual assumptions [are] needed to introduce an intended referent which has not itself been explicitly mentioned" and cases where the intended referent is mentioned under a different description in a previous utterance, as in (3) and (4) below (both of which are bona fide bridging inferences according to most accounts):

(3) I met a man yesterday. The nasty fellow stole all my money.

(4) Peter took a cello from the case. The instrument was originally played by his grandfather.

Her definition allows for cases of bridging that are not normally considered by current theories, where the bridge is not to previous inferences but to salient background assumptions as in (5):

(5) [Peter and Mary are off to visit a flat]. Mary: I hope the bathroom is not too small.

Is Matsui right in assuming that the cognitive tasks involved in fixing reference in (2) and (5) have more in common than either does with the task involved in (3) and (4)? The issue is of obvious psycholinguistic relevance.

Contrary to accounts that rely, for constructing the bridge, on the explicit linguistic information in a prior utterance (Clark, 1977) or on a situational model (Garrod and Sanford, 1988; Walker & Yekovich, 1987) , Matsui predicts that "in interpreting an utterance, the individual automatically aims at optimal relevance [which means] he will try to pick out, from whatever source, a context in which to process the utterance so that it gives at least adequate cognitive effects for no unjustifiable processing effort." This prediction is supported by a series of investigations based on utterances (presented alone or in the context of a story), with two plausible intended referents. Consider (6):

(6) I prefer the restaurant on the corner to the student canteen. The cappuccino is less expensive.

Is it the cappuccino at the restaurant or at the canteen that is said to be less expensive? 80% of Matsui's participants indicate that one can generally get less expensive cappuccinos at student canteens. If such common-knowledge were the determining factor, then participants should construct a bridge from cappuccino to student canteen. Similarly, if the determining factor were the shortness of the gap between the referring expression and a previous expression to which it could plausibly be bridged, then the canteen, the mention of which is the closest to that of cappucino, should provide the preferred bridge. Yet 100% of participants respond restaurant when asked "Where is the cappuccino less expensive?" Unlike theories developed in psycholinguistics, relevance theory provides an explanation of these data. The sentence The cappuccino is less expensive achieves optimal relevance as an explanation of the speaker's preference for the restaurant over the canteen when the bridge is to the restaurant, and is of no obvious relevance if the bridge is to the canteen. This is why all participants understand the phrase "the capuccino" to refer to the cappucino at the restaurant. Matsui's work provides striking examples of the mutual relevance of pragmatics and psycholinguistics (for further discussion, see Wilson & Matsui 2000).

3. Pragmatics and the experimental psychology of reasoning

Fruitful interactions between pragmatics and experimental psychology are not limited to psycholinguistics. All experiments involving verbal communication with participants are affected by the way in which they understand what they are told. When an experimenter's expectations do not measure up with a participant's comprehension, this can have major consequences. In the psychology of reasoning in particular, experiments typically involve not one but two levels of verbal communication from experimenter to participants: verbal instructions on how to perform the task and the task itself consist partly or wholly of verbal material. Experimenters (who are usually focused on rates of correct responses) often take it for granted that instructions and the verbal material are understood as intended, but this need not be the case. What happens if the instructions or text for a reasoning problem are not understood as intended? The performance of participants may fail to meet the experimenters' criteria of success because they have, in fact, performed a task different from the one intended. Their pragmatic comprehension processes may be functioning quite properly, and so may their reasoning processes, and yet their responses may seem mistaken to the experimenter. This is enough to give some plausibility to the claim that participants' apparent irrationality in reasoning tasks is linked to miscontruals or reconstruals of the task rather than to their reasoning incompetencies (Henle, 1962) . Even apparently successful performance of a task may in some cases be due to an unforeseen interpretation that happens to yield the experimenter's normative response not for logical but for pragmatic reasons.

The role of pragmatic processes in reasoning experiments is generally acknowledged, but only in a vague sort of way. There has been no attempt to introduce systematic pragmatic considerations into experimental methodology. Nevertheless, there have been more and more studies investigating the role of pragmatic factors in standard paradigms in the psychology of reasoning, following the pioneering work of researchers such as Politzer (1986) and Mosconi (1990). A number of apparent irrationalities in people's performance have been shown to be explainable, at least in part, as resulting from these pragmatic factors. It is not an exaggeration to say that nearly every task in the reasoning literature has inspired a pragmatic analysis. Several illustrations can be found in the chapters by Politzer and by Van der Henst and

Sperber. The relevance of this work to the study of reasoning is self-evident. Its relevance to the experimental study of pragmatics is also clear because, in each case, researchers have had to identify precise pragmatic factors at work and devise ways of testing their role experimentally.