

Preface

The argument in this book aims to apply a body of cohesive and interpretable ideas, developed over the last dozen years or so, to issues of significance in educational psychology and epistemology. The history and development of these ideas, which emerged from experiments on perceptual motor learning, group interaction and sequential choice (as well as more obviously relevant studies of learning, subject matter structuring and cognition), are described in two previous books (Pask 1961, 1975a). But the main themes are crystallised in a monograph (henceforward called "the previous monograph"), Pask 1975b, *Conversation, Cognition and Learning*, which is part of the present series. In fact, the previous monograph marks a point of departure, for the notions cling together well enough to count as an empirically supportable theory: Conversation Theory.

Ideally, perhaps, *Conversation, Cognition and Learning* should be read first. But there are some 600 odd pages of it, including some lengthy appendices, and provided the reader will take various statements on trust, it is quite possible to start with this book. *Conversation, Cognition and Learning* can be regarded, with equal legitimacy, as an essay in man/man and man/machine symbiosis or as an essay upon education, learning and the like. In contrast, the present book is an application study and is unambiguously oriented towards the areas of education, its psychology and epistemology. The Introduction provides the essential groundwork, and for those who have read *Conversation, Cognition and Learning*, it bridges the gap between the two volumes.

Technical jargon has been minimised and examples have been

stressed in order to increase readability. But it is also true that a good deal of fresh ground is broken. There has been progress both in the theoretical and empirical areas since 1973, and the picture which can now be drawn is more readily comprehended and rather more comprehensive. The theoretical and experimental work is focussed upon learning strategy and style, upon innovation and "learning to learn," and upon the representation of knowledge by teachers, students or subject matter experts. The enquiries in both areas lead to some novel perspectives and discoveries.

Right at the outset I would like to qualify this pretentious word "discovery". One of the lessons continually relearned by our research group is that most of the "discoveries" amount to a restatement (with suitable backup) of the intuitions and covert opinions entertained by well-informed educators; so their surprise value is less than it might be. For example, styles and strategies of learning and problem solving are known to exist: Understanding is often conceived as *some* kind of reproductive process. All we do, in this respect, is to assert that there are particular kinds of strategy and that an understanding is a particular type of reproduction; that is, to render the common belief explicit. The position is a little different in the epistemological arena; some of the comments upon the nature of knowledge are surprising and uncover an interesting cognitive pattern. Moreover, the *methods* used both for subject matter structuring and the detailed study of individual or group learning are (I think) genuinely novel and merit attention as candidates for general employment.

Another lesson we continually relearn is that originality is something of a snare if not a positive delusion. Other people have thought the same thoughts and sometimes done the same things while using different idioms and methods, which frequently obfuscate the unmistakable similarities. Some debts and dependencies were picked up in the previous monograph; in this book there is a determined and fairly systematic attempt to establish the proper linkages and set the work in the context of the entire field.

Sometimes this is a difficult task. Commonly enough one is unaware of an intellectual debt except in retrospect and this is especially true when the donor speaks from a different platform. For example, all system theoretic and information process oriented psychologists owe an immense amount to Craik, working chiefly with Bartlett; but the magnitude of this particular heritage only

became evident when Craik's notes, essays and memoranda were edited for publication by Sherwood (Craik, 1966). By the same token most people concerned with knowledge and its representation have (often quite unconsciously) garnered ideas from Merrell (1966). It is clear, at any rate, that many of my own "original ideas" recapitulate the argument in his *Epistemics* and were probably born in discussion with the author some 15 years ago. Much the same comments apply to Grey Walter. Nearly everything worth saying that is said in this book about concurrency and local synchronicity (of a priori asynchronous systems) is contained in a prescient article (Walter 1956; the paper was presented in 1953) where the mechanisms in question underlie a phenomenon happily named "Abcission". Moreover, transplanted from cognitive studies into neurophysiology, the experimental methods he devised for displaying and quantifying this phenomenon are virtually identical with our own methods.

Other acknowledgements are quite easily tracked down. The likeness of conversation theory to the theoretical underpinning of the Vygotsky-Luria school and the Piaget school was evident from the outset but only became obtrusive after lengthy and illuminating discussions with Michael Huberman. Chapter 1 is mostly concerned with bringing the pertinent methods and techniques into register with the standard experimental conditions appropriate to conversation theory.

Substantial portions of the book were rewritten after a series of seminars and discussions with Gergely (a collaborator of Ivanhanko) and Nemeti occasioned by their recent visit to Great Britain; clearly, they and their colleagues are saying the same things (more elegantly from a mathematical point of view) insofar as they have consistently applied their concepts to social systems and the development of science and have pursued their research over more than a decade. In view of this fact it would plainly be impertinent to construct an ad hoc "string and sealing wax" calculus to replace well-developed notions. In using these notions, in grossly simplified form, as a cornerstone of the argument I hope I have neither misrepresented their position nor distorted a very beautiful theory. Their own books on the subject are in preparation.

I am particularly indebted to John Daniel both for helpful criticism and inspiring ideas; for example, that entire educational sys-

tems can be characterised on a par with individual or group learning (his remarks on the divergence between educational styles, Daniel 1974, is well worth consulting). Marvin Minsky's theory of "Frames" turns out, on discussion, to be more or less the same as our theory of concepts. Nicholas Negroponte, in many ways, is responsible not only for the basic ideation but for extending it to the wider horizon of design and architecture — quite apart from his role as tutor in how to implement man/machine interaction.

In the previous monograph I stressed the conjoint origins of much of this research and noted that it stemmed from the intellectual mandates of McCulloch and Ashby. That is still true. Most of the research lines have also been pursued simultaneously but more or less independently by Von Foerster and his collaborators (Loefgren, Gunther, Weston, for example). Over and above these dependencies which were mentioned in the previous monograph, this book owes a great deal to the fresh efforts of other colleagues. As often in the past, Prof. Brian Lewis and others at the Open University have commented upon and criticised the manuscript; Lewis has read it in detail and his revisions have been freely incorporated. I owe a debt to my associates at System Research (to the extent that this is really a compound document): in particular Robin Bailey, David Ensor, Dionysious Kallikourdis, Robert Newton, Elizabeth Pask, Valery Robinson, Bernard Scott, and Tony Watts. Most of the ideas have been refined and several of them instigated by faculty members (especially Laurie Thomas) and graduate students at Brunel University and at the University of Illinois at Chicago Circle (where a couple of chapters were written). There I am specially indebted to Professors Conant, De Fanti, Laxpati, Bruce McCormick, Manacher; to Ted Nelson and others in the Department of Information Engineering; to participants in the Applied Epistemology Seminars; for example, Sally Drogue, Professor McNeil, Professor Miller, Dr. Joe Lipson, the Tiemans, and to Larry Leske, Dave Douglas, and Randy Walser. Laxpati, Leske, Douglas and Walser made detailed and valuable criticisms of the manuscript; in addition they have realised an implementation of several operating systems within a slightly modified form of PLATO.

I would like to thank Isaac Haissman of System Research for scientific editing and the preparation of an index.

By a conventional impropriety the most important people come

last. Our research group is a Social Science Research Council Research Programme: "Learning Styles, Educational Strategies and Representations of Knowledge: Methods and Applications," and the research is carried out at System Research Ltd. For the most part this book is an account of this programme, its ambitions and its achievements (occasionally it goes beyond our brief though not, perhaps, our endeavours). These patrons are not only sponsors but valued advisors.

Linda Barsby has typed manuscript drafts repeatedly, corrected them and often the author. Bernard Scott and Robin Bailey have read it and Scott is responsible for the detailed referencing.

Gordon Pask

Introduction

The previous monograph (*Conversation, Cognition and Learning*, Elsevier, 1975) deals with the history and implementation of techniques designed to exteriorise cognitive operations, especially those of learning and of teaching, so that they can be observed as segments of dialogue and behaviour. One method of exteriorising cognition is to engage in a verbal conversation, with a learner for example, and to *discuss* the way he learns as he learns.

This method has several obvious defects. The dialogue interferes with progress. The experimenter loses his status as an external observer, since he participates in and biases the learning process. Natural language expressions are hard to interpret and may be inherently ambiguous. Even so, the amount of information about mental events which can be obtained by this means greatly exceeds the amount obtainable by the classical type of stimulus/response or input/output experiment. In fact we proposed that as the classical type of experiment is improved to approximate the ideal (the respondent is isolated in controlled and replicable conditions), the information available to an external observer regarding conscious operations will decrease very rapidly to the vanishing point. Conversely, the information about conscious operations is maximised by establishing an appropriate kind of dialogue which is overlooked by an external observer.

Some of the difficulties mentioned in the last paragraph can be surmounted. For example, it is possible to distinguish the roles of external observer and participant experimenter; the observer gives instructions to a participating agent (how to act and what to discuss), after which he looks on dispassionately. The agent in ques-

tion may, for many purposes, be either a human interviewer or a mechanised system. Much of the argument in the previous monograph was couched in terms of a mechanised system (CASTE: "Course Assembly System and Tutorial Environment"): partly to make a clear distinction between what can and cannot be mechanised and partly as a practical expedient (human beings are unable to sustain the role of participant experimenter if the conversation ranges over sizable subject matter areas, if the "instructions to the agent" are precisely obeyed and if transactions are to be recorded).

To meet another objection, the conversational language need not (for many purposes) be natural language. The conversational language (henceforth designated L to distinguish it from the meta-language L*, employed by an external observer to *talk about* the conversation) may be a graphic or non-verbal symbolic language. Certainly, L must be quite a rich language. For instance it must be a programming language as well as a descriptive or assertoric language; there are genuine L questions and L commands (not just formal surrogates for questions and commands); L statements must refer to *persons* "I" and "you" as well as objects; L must have an unusually liberal interpretation or *semantic*. Even so, it is often possible to realise the L transactions as sequences of concrete operations and in that case to replace verbal utterances by behaviours which can readily be computer monitored and recorded.

Some caution is needed when using the word "behaviour" in this context. The necessary caveats were stressed in the previous monograph to produce an almost obsessive notation in which behavioural terms like "stimulus" and "response" were generally eschewed. Having made the point, it is legitimate to relax the nomenclature provided that the behaviours attending L transactions are recognised as *many sorted*. (In contrast, the most extreme forms of behaviourism view behaviours as *one sorted*; a precondition for synthesising complex entities out of simple ones, or conversely, for an atomistic analysis of complex behavioural events.) To illustrate the *many sortedness* of behaviour we should distinguish between simple behaviours (causally, albeit probabilistic-causally determined) which are the *one sort* of classical behaviours; model-building or rule delineating behaviours (a sort of behaviour that delineates an explanation or a demonstration); and *learning strategy* behaviours which represent, by a concrete tracing, how an

explanation is derived from other explanations.

Depending upon the form of L, there are many types of dialogue which will exteriorise mental events and they are graded or typed in a series extending from free natural language dialogue, via restricted natural language dialogue, to situations in which L transactions are mechanised. All of these types are called *conversations*; the necessary experimental methods are called "conversational techniques".

The objection which cannot be eliminated, whichever technique is used, is that any conversation takes place within a contractual or normative framework. The respondent *agrees* to engage in the conversation, for example, in order to learn about a subject matter, and this agreement or contract is *negotiated* in natural language L*, though it may also be *expressed* in L. Further, the participating agent, either man or machine, biases the conversation: literally an external observer looks at a *conversation* not at unfettered response (whatever that may be). In aggregate, these objections are not very serious. The price paid for observation is no greater than the price paid in a classical experiment though the biases and constraints are manifested differently. Moreover, at least in systems like CASTE, the amount and type of bias can be estimated *after the event*, though it cannot be accurately predicted beforehand.

We now come to the underpinning contention of the previous monograph. Psychological phenomena, especially those involved in learning and education, stem from or are related to states of *consciousness*. Using the argument which relates the information available about conscious processes to the type of experimental situation, we maintain that the basic unit of psychological/educational observation is a *conversation*. In order to test hypotheses and explicate the conversational transactions, it is necessary to invoke various tools and explanatory constructs. These are coherent enough to count when interlocked as a theory, and this theory was dubbed *conversation theory*.

1. PREREQUISITES

Certain prerequisites are demanded of any worthwhile theory.

1.1. Observation

To begin with, there must be at least one *standard condition* for measurement and observation; other conditions, usually less restrictive, being systematically derived from it. The standard condition of *conversation theory* is called a *strict conversation* and it is possible to instruct or program participating agents so that if *any conversation takes place*, then it is certainly a *strict conversation*. The main features of a strict conversation are as follows:

(a) The participants, as part of a contract, agree to obey the rules of the conversational language L and the participating agent makes sure that the L syntax is respected.

(b) The conversation is focussed, or anchored, upon a conversational domain: typically, a representation of the topics in a subject matter.

(c) The conversational domain involves a particular and canonical type of representation, both of what may be known and what may be done or discussed: hence, conversation theory has an epistemological commitment, and about half of this book is devoted to an exploration of what this commitment is.

(d) Each topic, said to be learned or assimilated in a strict conversation, is understood.

(e) In this connection *understanding* is given a theory specific and technical connotation though the imputed meaning tallies with and probably amplifies the usual meaning. We say that a topic T is understood by a participant if and only if T is explained and if T can also be derived from other topics in the conversational domain, i.e., a derivation is an explanation or systematic justification, *of* an explanation. It is crucial that understandings can be detected.

The explanation need not be verbal. If not, then it is called a model-building operation and is a satisfactory explanation insofar as the model can be executed in an external facility to bring about the formal relation underlying topic T. Nor need the derivation be verbal. If not, it is a *learning strategy* (a concrete depiction of one or more topic derivations).

(f) A strict conversation is punctuated by *understandings* and the intervals occupied in reaching an understanding are called occasions.

1.2. Framework admitting inference

Another prerequisite for admissibility is that a theory shall have predictive power and that its predictions can be empirically falsified when tested under the standard condition. The predictive capabilities of conversation theory chiefly emerge from psychological or systemic postulates introduced in order to furnish a mechanism of understanding.

The critical mechanism-postulates developed in the previous monograph are as follows: Concepts and memories are regarded as dynamic constituents of the mind. Specifically a concept is regarded as a procedure that realises or satisfies a topic and the topic itself is an interpreted (formal) relation. For generality, we say that a concept is a procedure that reconstructs or reproduces a topic (T). By virtue of this definition it is natural and in line with ordinary language usage to assert that a memory is a procedure that reconstructs or reproduces a concept. We contend that stable concepts, for all practical purposes the concepts existing in a mental repertoire, are those which can be reconstructed or reproduced by at least one (usually many) memory-procedures in the same repertoire. It follows that learning is an evolutionary type of process in which concepts and memories are constructed, *ab initio*, and an understanding signifies the generation and existence of a stable concept, i.e., a concept associated with a memory which either exists or is created in the process.

These definitions fit in quite neatly with the events observed in the conduct of a strict conversation (which is not surprising since the postulates were advanced as plausible and worthy of serious consideration just because certain kinds of dialogue can be observed). Notably, if we looked at the execution of a concept inside some processor (programmable computing system) such as a brain, then the reproduction of a concept would appear as a cycle or series of execution steps and the instructions making up the procedure as a "listing," i.e., a series of linked statements which specify the intention or rule (in the same way that a computer program specifies a rule which, on execution, performs a computation). Similarly, if we looked inside the processor, the reconstruction of a concept (by a memory) will be manifest as a cycle or series of execution steps and the memory itself, as a series-like "listing". All this is a straightforward consequence of regarding concepts and

memories as replicative and reconstructive operations which can be described, in the abstract, in terms of several theories of self-reproducing automata. (Much of the previous monograph was concerned with the hedges and conditions needed to fit existing versions of self-reproducing automaton theory to mental reality; for example, mental operations are not generally serial or completely synchronous.)

Suppose that the cycles which might be examined by probing inside a brain, qua processor (call it α) are literally pulled out so that some of the cycle is executed in α as before and some of it in a distinct brain or processor called β . Under these circumstances, scrutiny of the interaction between α and β will expose the cycles to view. In particular, there will be one cycle corresponding to the execution of a concept, one to the listing of the concept and others corresponding to the listing and the execution of a memory. These are identified with stretches of dialogue or behaviour, as follows: the execution of a concept (to realise or satisfy a topic) is an exemplification (dialogue term) or a simple behaviour; the listing of a concept is an explanation (dialogue term) or a model-building operation (behavioural term); the memory cycle is a derivation (dialogue term) or a *learning strategy* (behavioural term).

A strict conversation gives rise to a series of transactions that are characterised as occasions, insofar as each topic learned is associated with an *understanding* (in its technical sense; a linguistic event involving the explanation and derivation of a topic or the construction of a satisfactory model for a topic within the framework of a learning strategy). α and β figure as the brain of a participating respondent and either the brain of a participant experimenter or a suitably programmed mechanical agent. The observable event of understanding is held to signify or evidence the construction of a stable concept due to a very specific kind of cooperative interaction between the conversing participants.

The circumstances under which cycles of explanation and derivation may be "pulled out" or (equisignificantly) "exteriorised for external observation" are precisely those set up by the contract and conduct of a strict conversation. In particular the learning participant must have a need to cooperate (implicitly identified with "procedure sharing" or "program sharing") in order to learn the topics in a conversational domain, which he has agreed to

do in the initial contract. The other participant must be in a position to provide this cooperation and foster understanding. Finally, insofar as "procedure sharing" or "program sharing" depends upon local synchronisation of the brains or processors involved, the occasions of a strict conversation are intervals of partial synchronisation between the participants during which they both attend to the same topic. Notably, such occasions are rare in nature. Brains, unlike computing machines, are not a priori synchronised by a master clock and it takes an act of attention (one type of "provocative transaction" noted in the previous monograph) to secure synchronicity.

A satisfying relation is thus established between the dialogue or behaviour of a strict conversation and a fundamental notion of information-transfer, due to Petri (1964). According to Petri's argument, information-transfer takes place if and only if two or more dynamic systems which do not have a common master clock (i.e., a priori asynchronous systems) come into local synchronicity (local with respect to topics in the conversational domain). Under this interpretation, the *occasions* of a strict conversation are, as they intuitively should be, indicators of information transfer.

1.3. Observable Units

We pointed out, in the previous monograph, that a theoretical framework of this kind permits several alternative definitions of a participant. Which definition is adopted is, to some extent, a matter of elegance and convenience.

Two perfectly valid alternatives are as follows:

(a) A participant is identified with a brain able to act as a processor for L-Procedures (henceforward, an L-Processor). The brains (L-Processors; α and β) are spatio-temporally demarcated on biological or mechanical grounds, the usual criteria for isolating an integral object in the environment. If participants are identified in this way, they are mechanically individuated (for brevity M-Individuated) by the external observer and count as Mechanical Individuals or α -Individuals. * By the same token other parts of the

* For most purposes "biological Individuals" would be just as acceptable. However, the class of L-Processors is larger than the class of brains, and conversely, brains have functions other than L-Processing.

environment (usually having less computational versatility than an L-Processor) can be M-Individuated, for example, various components of CASTE or any other design of experimental situation.

(b) A participant can be identified with the set of stable concepts that are, or may be, part of his mental repertoire. To obtain this characterisation of a participant it is only necessary to extend the sequence of formulations "concept, memory, ..." until it is possible to answer the question, "what reproduces the memories that stabilise the concepts, thus yielding a unitary and recognisable repertoire." Since the answer to this question consists in a series of interlocking and compatible L-Procedures that are executed to realise a system of coherent beliefs or hypotheses, we say that the external observer has psychologically individuated (for brevity P-Individuated) the participant. If the constituent procedures are actually executed in some L-Processor, the participant is characterised as a Psychological Individual, or as a P-Individual. Although there must be some L-Processor to realise a P-Individual, we need not dogmatise about which processor it is (α , say, or β) and it often turns out to be impossible to do so. In this sense, P-Individuation is "processor independent". For example, the strict conversation is a P-Individual and is the direct object of observation. The participants, call them A and B, form the factor P-Individuals (A, B) of the conversational P-individual. Clearly, the execution of the conversational P-Individual is distributed (by "procedure sharing") over the M-Individuated processors α , β and its factors may be. This identification scheme also accommodates such obvious and important internal (and *not* directly observable) conversations as "thinking to oneself" or "learning on one's own account" (the coexistent execution of A and B in one brain, α) and group learning (where A, for example, is distributed over several brains α , β). As hinted already α , β need not necessarily even be brains (there are some inanimate L-Processors).

Although either formulation is legitimate the P-Individual is usually a more convenient unit for conversation theory; for example, a strict conversation is a prototypical P-Individual and the use of this formulation avoids a number of puzzling pseudo questions like "where did the concept come from?" or "which brain does it belong to?"

1.4. *Changed Emphasis in the notation*

In this book we do not make much explicit use of M-Individuation though such an act is implicit whenever brains are considered as distinct and recognisable entities. Moreover, the discussion often rests upon entities that are identified by M-Individuation: notably L-Processors (human brains and certain inanimate systems) and modelling facilities. The latter are vehicles in which models are manufactured as non-verbal explanations and (facilities of a distinct kind) in which derivations are reified as learning strategies. All modelling facilities have a dynamic component, they are computers and execute the models built in them to realise or satisfy relations: however, they are much more restricted computers than a human brain.

P-Individuation is, however, used quite extensively and the P-Individual, as a unit, is ubiquitous. As in the previous monograph, a P-Individual is realised by execution in an L-Processor and, generally, one or more L-Processors are assumed to be available. If that is not the case, we distinguish (notably in Chapter 11) between the representation of P-Individuals A, B written Π_A, Π_B and the P-Individuals undergoing execution (just A, B, *simpliciter*).

Since a great deal of the argument is concerned with the creation and learning of analogy relations of a much more general and useful kind than those discussed in the previous monograph, we often need to pay special attention to the interpretation of a topic (its realisation in some universe, in contrast to the formal systemic or syntactic topic relation). This trend penetrates to all levels of the argument and motivates a change in notation, though not in principle, from the standards established in the previous monograph. In order to treat interpretations and analogies intelligibly, it is desirable to discriminate between programs as syntactic entities, the compilation of programs (the configuration set up in an L-Processor or any other computer, which is open to execution), and the execution of the compiled program.

Whereas before, programs written in a modelling facility (as non-verbal explanations) were identified piecemeal with models as "compiled programs, compiled in the modelling facility," it is now more expeditious to distinguish the program (i.e., the listing as a syntactic entity, to call "the program compiled in a modelling facility" a *model*, and to consider the execution of the model.

This usage accords with recent theories of semantic interpretation that are usually called "model theory". *

Under these circumstances, the notation *Exec* (used in the previous volume to designate the execution of a procedure in order to produce a program *listing* which was compiled in an external modelling facility to produce a model) is positively misleading and is herewith discarded. Its main virtue, in any case, was to complement the explanatory response *Expl*. Throughout this book (unlike the previous monograph) we speak explicitly of non-verbal explanations as the production of program listings which represent the class of programs making up a concept. We refer to the compilation of such representative programs in a modelling facility as *models* and, when necessary, refer to the execution of models in a modelling facility (under the control of the facility and not under cognitive control). This brings the argument back into kilter with the previous monograph. But the revision admits a relatively uncomplicated account of analogy relations and their models (the non-verbal explanation of analogies), a topic which often dominates the present discussion.

Corresponding notational adjustments are required in respect of concepts, memories and P-Individuals. Whereas, in the previous monograph, these were regarded piecemeal as procedures under execution in an L-processor, it is now expedient to discriminate a syntactic component of each entity (its formal specification) which is called a program, a Procedure being a compiled program. Hence a concept is respecified as the stable compilation of a program in a brain or other L-Processor; a memory as the stable compilation of a different kind of program, and with one caveat, the P-Individual as the stable compilation of properly adjoined programs. Hence, concepts may be selectively executed provided they are stabilised (as compilations) by memories. Similarly, memories may be selectively executed to stabilise concepts. For the P-Individual there is an additional requirement; namely, that some of its programs (compiled as concepts and memories, in general, as procedures) are invariably undergoing execution.

* Until recently model theory was mostly concerned with static models. In contrast the current argument is almost exclusively focussed upon dynamic and executable models (i.e., compilations of programs of one kind or another either in brains or mechanical artifacts).

We refer, throughout, to L-Processors (in which concepts, memories and P-Individuals are executed as programs and procedures). This usage may *seem* to be (in fact it may *be*) eccentric since an L-Processor is nearly always a brain. But on balance it has value insofar as it does bring home the following facts: (a) An L-Processor may be a brain or a collection of brains or a man/machine system, without prejudice; (b) Not all of the brain can act as an L-Processor and a brain has other functions to perform.

1.5. Testability of Postulates

Conversation Theory should have predictive power and be open to falsification under its standard condition, the strict conversation. It indubitably does have predictive power and its predictions are open to falsification. For example, we predict that the concepts of understood topics shall be indelible within one conversation if it is anchored upon one conversational domain, and that they should be relatively resilient to the interference effects encountered if they are recalled or executed in a different and perhaps incompatible conversational domain. We also predict the existence of classes of learning strategies which become mutually exclusive in a strict conversation, for example, the previous monograph emphasised holist learning strategies and serialist learning strategies. A fair body of empirical evidence, supporting these and more subtle hypotheses, is collected in the present book, and the tenure of hypotheses in conditions that deviate from the standard conditions is examined at some length because many educational situations *do* deviate quite markedly.

On the whole the salient hypotheses are supported by the experimental data and many of the tenets of conversation theory continue to hold (sometimes with modification) under circumstances that are very realistic (and often very deviant). It is wise, however, to stress the status of conversation theory and to consider what it can and cannot be expected to do.

1.6. The Scope of Systemic Theories

As outlined so far conversation theory is a systemic microtheory or molecular theory. It refers to concepts, memories and the like manifest in detailed transactions: either stretches of dialogue, or

stretches of many faceted behaviour. Although the theory is detailed and mechanism-oriented, the mechanics are systemic (i.e., patterns of organisation) and the theory is not intended to discriminate particular biological processes (for example, any or none of the very different memory mechanisms proposed by Bogoch, John and Ungar and discussed by Libassi [1974] may be responsible for the compilation of our "memories" or our "concepts"); the theory is neutral on this score.

To some extent, this degree of neutrality is maintained with respect of functional distinctions as well. To illustrate the point, consider the learning theory proposed by Atkinson and Shiffrin (1965, 1967). This theory provides a bridge between statistical learning theories and the "artificial Intelligence" approach to mental activity (for example, Feigenbaum 1959, Feigenbaum and Simon 1962, or Simon and Feigenbaum 1964). It posits a structural demarcation of storage media; a sensory buffer, a short-term store and a long-term store, with appropriate connections. Although these storage locations have specific properties and capacities they are functional loci; not (except by indirect inference using other evidence) sites in a brain. Within these locations and the constraints they impose, control processes, which are identically "compiled programs or procedures, undergoing execution" set up and manipulate data structures — for example, a rehearsal buffer is maintained in short-term store and other control processes, which generally compete with rehearsal buffer operation, select symbols for acceptance into the rehearsal buffer.

Of course, this is also a *systemic* theory. However, its validity (there is strong evidence that it provides an excellent picture of short-term storage, at any rate) neither confirms nor denies our theory, or vice versa. Notably there *might* be competition; it simply happens that some mechanism of this type is mooted as part of our own theory and almost any mechanism would do. Moreover, the detailed transfer patterns (between structurally demarcated compartments) are represented statistically as a result of which the content of the theory neither confirms nor denies the kind of cyclic reconstruction we posit (though the rehearsal buffer process could surely be regarded as *one example* of a cyclic reconstruction process).

1.7. Comparison between conversation theory and other systemic theories

In the previous monograph, Chapter 11, we noted that a macro-theory or molar level conversation theory is possible and some effort was made to relate subjective uncertainties (sampled by confidence estimates and the like) to the activity of mental systems. The macrotheoretic variables are, or are derived from, degrees of doubt and certainty. We distinguished in particular, certain *kinds* of doubt: d_4 or doubt about what topic is being attended to; d_1 or retrospective doubt, given that a topic is understood, of which one of several methods (all belonging to the topic's concept) *is used* to solve problems posed under the topic on a particular occasion; and d_2 , or prospective doubt, given that a topic is in the field of attention but is not understood, about alternative outcomes or solutions to be obtained by applying the existing, and perhaps partially formed, concept. Moreover, we specified a "look ahead" uncertainty; namely, a doubt, given that *one* learning strategy *must* be selected from a set of possibilities, about which one will be selected.

All of these quantities are measurable, and from time to time, we take advantage of this fact. Also, at the macrolevel, conversation theoretic predictions do mingle with the predictions of other information processing and systemic theories whenever the experiments are comparable. So far as we can see (and there is, as yet, rather little experimental overlap) our own results are in accord with those of other researchers. This is especially true in the context of recent results on the perceptual and cognitive psychology of recall and recognition, a body of data far richer than our own limited scope experiments. Though I have not attempted to do so in this book, it appears that our findings can be transformed, by change of idiom and context, into substantial agreement with these results (e.g., the *Attention and Performance* publications).

Results from experiments in conversation theory, as the theory stands at the moment, are directly comparable with results from information processing theories and the psychology of "Decision Formulation" (that is "Decision Making" insofar as it refers to heuristics or mental operations, rather than the art of weighing up alternatives).

Some representative information processing theories are those of Broadbent's later work (in and after Broadbent 1971), or of

Cohen (1964, 1972) on subjective probability and choice tactics; theories of cognitive mechanisms (for example, Conrad 1974); the work of Daniel (1974) or Dirkzwager (1974) and his group (both the latter include replications).

It is quite possible that conversation theory can be developed to yield predictions/data compatible with the psychology of more elementary information processes; underlying the kind of cognition/behaviour in Welford's (1968) summary of the field, earlier with Broadbent and the Cambridge Applied Psychology Unit, or even the "signal in noise" treatment of perception and recognition pioneered by Tanner and Swets (1954). In order to bridge the gap between complex phenomena such as understanding and elementary mass phenomena (signal detection in a noisy background), it is necessary to provide a statistical treatment of memories, concepts, etc. This turns out to be a statistical mechanics (with some peculiarly psychological features) in which the dynamic systems making up the canonical ensemble are P-Individuals. On interpretation, the members of the ensemble may either represent students in a class (when the condition of the ensemble represents a state of general knowledge) or factor P-Individuals in one student (when the condition of the ensemble represents a state of knowing). Work in this direction has just started and parallels very closely the approach to the regulation of cellular metabolism adopted by Goodwin (1963). In Goodwin's case the dynamic systems are basic units involved in enzyme production (DNA, RNA, ribosomes and feedback from products produced by the action of the synthesised enzymes). Hence the equations of the dynamic systems are quite different and their often oscillatory interaction has a different form. But, in other respects we encounter very similar difficulties and insights. At least the approach is a workable and potentially useful way of viewing mental activity, and it is at this level that direct comparison between conversation theory and the statistically interpreted structural theories is logically sensible.

To illustrate "Decision Formulation," where complex mental operations, heuristics, and the like are in the foreground, we cite the work of Tversky and Kahneman (1971, 1973), of Philips (1973), or Edwards (1968). The only difficulty in comparing hypotheses or results is that "Decision Formulation" theorists generally concentrate upon the use and nature of heuristics, concepts, or whatever, whereas conversation theory is generally

focused upon their development.

Entwistle (1975) points out that quantitative information theoretic and decision theoretic methods could with advantage be employed in educational psychology, and it is clear from his paper that he means methods which are founded upon structural principles or mechanisms and consequently have a commitment to information processing. These methods are based on systemic theories (on a par with the examples just cited) and, although information measures are used as a common currency (as they are in any application), the methods are inherently more powerful than "information theory" used only as a metrical device. Entwistle's reasoning would (in our view, it *should*) find general acceptance. The trouble is that few relevant studies of this kind have yet been published though many of them are in progress.

Conversational domains (and, with them, the epistemological aspect of conversation theory) are also represented in systemic terms. Comparison with other work is relatively easy; in fact, an almost embarrassing number of comparisons are possible (many noted in the previous monograph and some to be introduced). For example, both data base design (at one extremity) and the semantic networks and data structures of cognition science (at the other) have features in common with our own formulation.

Probably the chief differences between conversation theory and other systemic theories are as follows: Conversation theory is explicitly relativistic; this is evident on inspecting the standard condition. Measurements are made relative to a conversation, or of one participant, relative to another, in the context of a conversational domain. Most of the other theories do not make the point explicitly, though some of them probably involve relativistic estimation. For the other outstanding point of difference, conversation theory is, with the possible exception of some events in a strict conversation, overtly reflective. It permits personalised statements "I" or "you" not just impersonal statements about objects and makes an attempt to explicate their nature. This, of course, is part and parcel of our general concern with consciousness as the distinctively psychological phenomenon.*

* The justification for relativistic and reflective theories is discussed at much greater length in Pask (1961) and in Pask (1975a), in particular, the development of pertinent experimental methods from studies of perceptual motor learning.

For psychology in general the merits of an orientation to conscious phenomena, to relativism and reflectivity are frequently debated. But whatever the outcome, it seems that a theory of this kind is required in order to deal with practical problems in educational psychology and the wider educational issues of course design, the structure of institutions and media, and the origin of creativity and innovation.

1.8. Unification

An incidental but valuable claim for conversation theory is that it unifies a number of psychological theories which otherwise appear entirely different. In the previous monograph, we examined several representative schools of thought in this light and tried to show the points of systemic identification between Personal Construct, Information Processing, Cognitive, Transactionalist, Behaviourist, and other psychologies by mapping them onto a conversation theoretic image. The present book goes a good deal further. On the one hand, the argument extends the domain of application from educational psychology to epistemology. On the other hand, the argument unifies various essentially conversational techniques (thus acknowledging the roots of conversation theory) and various theories of thinking, innovation, social learning and development.

2. A PLAN OF THE BOOK

Chapter 1 provides a survey of other conversational methods (Piaget, Vygotsky, Luria, for example). Although the present theory was developed independently (deliberate isolation in an attempt to integrate ramifying thoughts), it owes whatever value it has to precedents established in the culture and we try to trace the real origins, in retrospect. We also take the opportunity to describe the operating systems used in the experimental work: INTUITION (a transportable modification of CASTE, used in schools) and several others.

Chapter 2 very briefly reviews the structure of conversational domains as set out in the previous monograph, but most of the material is novel; we report work that has been done since 1973 to

yield an enriched and more generally useful product. In particular, the notion of an analogy relation is broadened (whilst the analogy is still represented systematically in a conversational domain). The significance of this manoeuvre is partly epistemological and partly practical. We posit that the rate of learning is materially influenced by the number (or density) of analogies a learner can appreciate, the quality of learning by the number of valid analogies that the learner comes to understand.

Chapter 3 reports a number of recent studies to do with learning strategies and styles; in the light of these results the holist/serialist distinction of the previous monograph is seen as an important but special case of more fundamental and pervasive mental processes.

Chapter 4 is concerned with theoretical developments bearing upon agreement and understanding and also upon the character and origin of analogies as "petrified agreements". The discussion in this chapter hinges upon independent work in two main fields; non-classical model theory and the coherence theory of truth. Both fields appear to be of the utmost importance to any rational theory of education—conversation theory or any other theory.

Chapter 5 furnishes a series of condensed notations or schemes for the description of learning. By adopting these notations, it is possible to avoid a great deal of symbolism (such as the symbols for complex transactions used in the previous monograph) whilst remaining in a position to describe the types of learning discussed in Chapter 3 and the acts of invention discussed later in the book.

Chapter 6 introduces the topic of conversations in which there are two or more simultaneous foci of attention, either on the part of several coupled participants (a group) or just one participant (a transient phenomenon, believed to underpin innovation).

Chapter 7 contains a description of a course assembly system, THOUGHTSTICKER, much more versatile than EXTEND (of the previous monograph) in which one or more subject matter experts maintain distinct foci of attention, from time to time, whilst building up a conversational domain.

Chapter 8 is also concerned with THOUGHTSTICKER but especially with transactions that lead to innovation.

Chapter 9 is devoted to an argument relating the art of course assembly as it is practised by experts (delineating knowables, constructing a conversational domain), to the art of "learning to

learn," as practised by students. We maintain that "learning to learn" is a crucial accomplishment and that a student who can do so effectively is (amongst other things) able to impose a personal structure upon otherwise unstructured information or upon an often perversely structured environment. Experimental data are cited to support this view.

Chapter 10 makes explicit a theory of creativity and innovation developed at various points in the preceding discussion and shows its relation to several other theories of innovation. It appears to tally with them all but is, in a systemic sense, more general (i.e., in this sense, it encompasses them as special cases suited to particular kinds of innovation).

Chapter 11 is speculative. It deals with work in progress and sometimes far from completion. But the issues addressed, such as characterisation or dramatisation, the nature of the media, the scope of developmental studies, strike me as fascinating and I hope the reader will find some of the novel perspectives both interesting and useful.

In conclusion, there is one general caveat. By disposition, I like to think as a philosopher (or a philosophical psychologist). To justify this mode of thought and to implement the conclusions experimentally, it is often helpful to build physical systems (INTUITION and THOUGHTSTICKER, for example). Under some conditions these are essential experimental tools, under other conditions they are valuable tutorial devices. Often, however, it is possible to realise the principles derived from experience with these systems in human terms, with human teachers in a classroom, subject matter experts working in a group, and in various other ways involving no machinery at all. On balance, we believe that most if not all of the findings and principles discussed in this book can be employed without invoking machinery (even though the discussion itself is machinery laden). Such non-mechanical implementations are usually of greater practical significance and may even be inherently more effective.