1 Introducing Cause

Cause and public issues

'Working Mums Blamed for Children's Failures' is a typical newspaper headline of today. As government and other organizations vie to shed responsibility, the supposed reasons for undesirable states of society are tossed around with abandon. These issues are so imbued with political, social and moral values that rational discussion seems impossible and not much changed over the past 20 years. How could working mothers be the cause of their children's failure? And what might one do about it if they were?

For the layman, there is a fairly obvious relationship between cause and remedy. So, if there is supposed to be a particular cause for an unwanted outcome, you can undo the outcome by removing the cause. It is obvious, isn't it? If the tap drips, fit a new washer. Although many government policies seem to be based on this principle, human problems are rarely that simple.[1]

Let us look more at working mums. The headline, quoted from *The Guardian*, was a very clear causal statement concerning the relationship between mothers working full time (as opposed to part time) and the school achievement of their children. This could be illustrated as in figure 1.1, where the arrow is intended to indicate a causal relationship, and the boxes are only there to package the text.

[1] A notable example is the logic that says that the way to stop an individual from committing another crime is to lock him up (this, of course, is a politician's logic, not a scientist's). Could it be, here, that it is the individual's liberty that is the causal factor? Take away the liberty and you take away the crime?

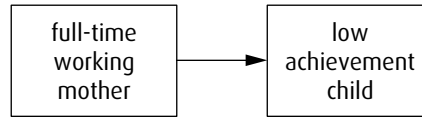

Figure 1.1

The remedy is clear. Use everyday reasoning: if A causes B, then if A is prevented, B will not occur. So, mothers with school-age children should be forbidden from working full time and this will result in increased school achievement! In these ways, strong government can be effective! Expressed like this, the conclusion is clearly monstrous. But why? The issue is not one of individual liberty. Suppose, instead of *compelling* mothers to work part time, that the state *rewarded* them for doing so. Would this be any better? This is one case in which common sense and other methods of data acquisition will be in agreement: it would depend. Without understanding the mechanism underlying the relation between mothers working and their children's school achievement, we cannot begin a rational approach to the solution. What is missing? One assumption is that mothers who move from full-time to part-time working will spend the surplus in the home and would thus be profitably available to the children.

Already, then, we see that the notion of full-time working as a cause is insufficient. We will return to the mother in a moment, but there are factors to do with the child which have to be addressed first. First of all, it must be apparent that, within the group of children in the study, there must be a vast number of individual factors that contribute to their performance. Let us for the moment include them all under the heading of the *child's state of mind* and represent that this is a factor in determining

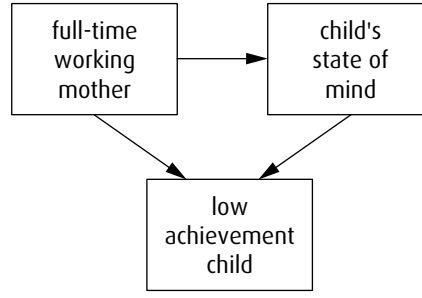

Figure 1.2

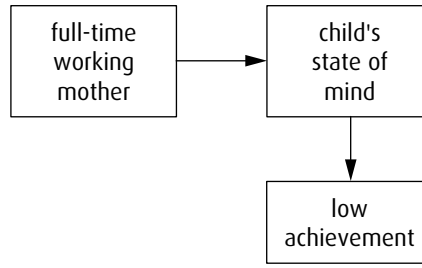

Figure 1.3

the level of achievement (behaviour) of the child. I diagram this in figure 1.2. This is not meant to be a profound thought, but we will see that it makes a difference to the way in which we think about the problem.

The next stage is to see that it is inappropriate to represent mothers working full time as having a direct effect on achievement. This is because the achievement referred to is a piece of behaviour – performance on school tests – which has an immediate cause, as it were, in terms of the intellectual capability of the child (plus other internal factors, such as motivation). The state of the child, then, has rather to be inserted into the chain of reasoning between the mother's behaviour and the child's achievement, as in figure 1.3.

In line with this, the causal diagram in figure 1.4 suggests that mothers who work full time are crucially not available to the children in certain ways. In other words, the absence of the mother causes something. What might that be? One suggestion is that if the children just returning home from school have nobody to talk to, they become

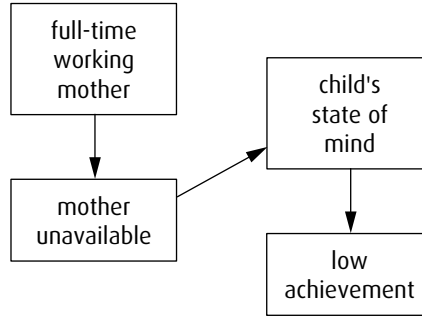

Figure 1.4

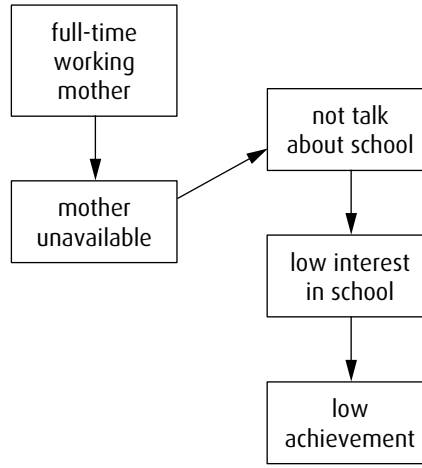

Figure 1.5

uninterested in school and this results in lower performance. This is diagrammed in figure 1.5. Such a statement would require some elaboration before it could be considered serious, and it might be interesting to consider the reverse relationship – that if the mother was home when the child finished school, the child would be able to talk to her and interest in school would be maintained. In other words, it seems that we still have a question about which way round the cause is; whether absent mothers depress performance or present mothers increase performance, or both. Before we can find a remedy, we have to discover which way round it is. If we were to go with the lack of interest in school as a cause, then we could correct it by finding some other way of increasing interest which did not involve the mother. But if it is the absence of the mother herself that is important, the lack of interest in school may be a symptom, and attempting to correct that (without the mother) would be pointless. Note that the data remain the same on all these interpretations, although the solutions differ.

Even with these moves, there are two obvious problems. The first is that there seem to be too many classes of exception, groups of mother/child pairs for whom the postulated relationship is manifestly incorrect. Examples would be children of depressed mothers, and those of mothers working full time at home. There is some research showing that the children of depressed mothers have low school achievement whether the mothers are at home or not. I leave it to you to imagine

the complications involved with mothers who are working full time at home. Note that *individual* exceptions in the level of performance are of no interest to this line of thinking. The thrust of the research is that certain effects obtain in general over the population as a whole. There will be both high and low performers wherever the mother is. For any individual, then, the question is not whether or not the child has done well or badly by some absolute measure, but whether he or she would have done better or worse if the mother had behaved differently. In general, such things are unknowable.

The second problem, already referred to, is that the mother, simply by being in the home, cannot be the *direct* cause of any change in intellectual capacity or motivation. We have to determine a *mechanism* to connect the two. Remember we are trying to understand the relationship between the mother's presence in the home and the child's educational achievements. *Being available* may be a useful concept, but is scarcely a mechanism. More specifically, one might argue that the mother's mere presence in the house at critical moments is not sufficient to produce the effects. Rather, the mother has to behave in particular ways, to be specified, in order to influence the child.[2] Even more likely would be an indirect effect, where the mother was fulfilling a pedagogical function, actually increasing the teaching time and the child's intellectual capacity. Such a proposition would be testable by looking at the mother's educational level, which would be expected to be closely related to her ability to teach the child. There are other possibilities, however, most obvious of which is that some mothers, perhaps by their presence in the house with the child and the interest that they show, have an effect on the child's motivation, which, in turn, increases the capacity to learn.

This extended, but still rather shallow, analysis illustrates how populist attempts to assign cause, blame and responsibility with the aim of correcting a problem are doomed to fail through their lack of subtlety. There are so many other factors, some of which could be major – for example, the effect of the absent father. In general, the scientific focus is missing. For a developmental scientist, the first trick is to define the problem properly. 'Children's Failures' is too broad a category to be subject to proper analysis. As we will see, the identification of the cause

[2] One mother I know of, expressing her concern over her 15-year-old son's tribulations, was told by him 'Why don't you find another interest?'

of a disorder is very much tied up with issues of diagnosis, treatment and management. But we will also see that classical development disorders, such as autism and dyslexia, although these are the terms used in diagnosis, turn out to be very complex when examined under the spotlight of causal analysis.

Cause and individual events: 'Why did Romeo die?' ■

Is the problem with working mothers that there are too many exceptions? Would we gain more understanding by looking at individual cases? To test this, I asked a number of friends why Romeo died. The responses were quite varied, but the general impression was that it was not a good question.

'OK, what *caused* Romeo's death?'

This was a little better. 'Love' was the most usual initial answer, but the question invariably provoked a lot of argument and successfully diverted people away from the war in the Balkans – the news topic of the day.

The kinds of answers that my friends came up with were varied and often strongly held. Was Romeo's death caused by Juliet's apparent death? By an underlying depressive state that ran in his family? By the drug Juliet took earlier, that persuaded Romeo that she was dead? By the feud between their families? Or by the parlous state of the postal service between Verona and Mantua?[3]

What we want to find are more or less proximal steps in the chain – cause and effect in a disciplined and systematic way. The most obvious answer is:

'Romeo was killed by the poison that he took.'

This is a scientifically acceptable answer. There is a gap in the causal chain – Did the poison lead to heart failure or brain failure? – but that doesn't matter if you know where the gap is and how to go about filling it. So, we might not know how that particular poison works

[3] In fact, Friar John, who had been entrusted with an explanatory letter from Juliet to Romeo, got caught up in a health scare and never managed to leave town.

other than that it acts quickly[4] and, with most actors, painlessly. But the kind of information that we would need to fill the gap in the chain will be clear. However, there is something rather pedantic about this answer.

Let us now take another answer that aspires to scientific status. Suppose it were the case that suicidal tendencies were heritable and we had evidence that the Montague family had some long history of assorted suicides, associated, say, with depression. We would still not be at all happy with a claim such as

'Romeo's genes caused him to commit suicide.'

Even if we accept that Romeo was depressive, it would have taken a remarkable conflation of circumstances to make him take the poison. A stronger personality would have said 'Mother was right. You can't trust those Capulets', but Romeo's genetic weakness showed up at that instant. Clearly, the jump from gene to behaviour in one go is too much. The gaps in the causal chain are chasms. One reason for being particularly unhappy with this explanation in the current scientific climate is that gene-to-behaviour statements give the illusion of having settled an issue, of having explained something, in spite of the explanatory chasm. On the contrary, I would claim that such statements only sketch one of the many jobs that have to be done. It should be clear that committing suicide is not one of things that genes code for.[5] So the job that has to be done will involve bridging the gap in some degree of detail. I insist on this for two reasons. The first is because it is becoming clear that almost every ability, trait or behavioural tendency is at least slightly heritable. In such a world, the claim of partial heritability for something, without some significant support to the causal chain, adds absolutely nothing. By a parallel argument, the bald claim that the environment exerts an influence on something is equally uninteresting. We have to be more specific.

The second reason for insisting on some kind of detail in the specification of genetic influence is that the probability of the outcome given the gene is so low, depending as it does on a multitude of environmental circumstances as well as on the presence of other genes.

[4] (For the benefit of all those people who believe that Romeo stabbed himself:) 'O true apothecary! Thy drugs are quick. Thus with a kiss I die.' *Dies.*

[5] Freudian notions concerning the universality of the death wish notwithstanding.

However, the specification of the environmental contribution to Romeo's death is equally shaky:

'Romeo committed suicide because he was brought up in a violent culture (where life was valued little and the means of killing easy to come by).'

Again, the gaps left by this sort of explanation are simply too great to allow the feeling of a satisfying answer. This is as vague an answer as the genetic explanation; the supposed cause here is much too ill-defined. Consider the intuitive psychological explanation:

'Romeo committed suicide because he thought that Juliet was dead.'

At first glance, this familiar phrasing in everyday language may look acceptable, but at second glance, it scarcely approaches the issue. This explanation, too, leaves enormous gaps. It implies that he did not want to live without her. However, if Romeo had thought that Juliet had died in a traffic accident, for example, on her way to see him, we could imagine that his response might have been different.

The drama, of course, does not use a single cause. Many factors conspire together to bring about the conclusion, which is all the more dramatic and poignant for having been multiply avoidable. Indeed, one might say that the whole play is a causal model for the finalé.

Everyday transactions provide a number of reasons for looking at cause. The most common, perhaps, is as a means of establishing responsibility for a particular event. In this way, we can establish blame (and our own innocence). Alternatively, we might want to know why something happened in order to find out what to do about the situation right now, or how to prevent the same thing happening again in the future. Such uses are not relevant to our current aims. Notions of responsibility or of cause of individual events are well suited to courts of law or detective novels. They are not usually suited to scientific questioning.

We are left, then, with some clues as to what is needed for a scientifically valid causal explanation. The example showed that cause and effect must neither be broad nor be too far removed from each other in conceptual space. In addition, it is clear that when we try to examine the individual case, it's easy to become overwhelmed with detail. There are a variety of individual factors that are central to the story and that

have claims for a role in determining the outcome in this individual case. Some of these factors are unique to this particular case, and knowing them would not help at all in understanding other cases or preventing future accidents. For other factors, there might be generalizations that could be formulated, such as *don't mess around with poisons* or *don't get involved with someone whose mother disapproves of you murderously*.[6] In the end, however, science has to deal with aggregates and probabilities – not with individuals and certainties – and a more appropriate framework for Romeo would be that of psychosocial pathways, where the primary concept is one of risk (of suicide) rather than cause (of death). Psychosocial pathways are touched on in chapter 3.

Some more reasons for not looking at individual cases

Suppose we were asked whether a particular eight-year-old would become delinquent in the future. We would need a variety of information, some to do with the child and some with the current environment. Perhaps we would ask first whether the eight-year-old was a boy. Knowing the sex of the child will get us a long way – we have information that violence is much more common in males – though it would be strange to attribute maleness as a *causal* factor.[7] Instead, being male is what we would call a *risk factor*, following Rutter (1989). Certain personality characteristics will also be important. Furthermore, we would want to obtain information on the child's parents: their past history, their employment, social class and marital status, the degree of marital discord and their current income. In this information, the contribution of social and genetic factors is unknown, and the interaction between them is extremely complex. Strictly speaking, in order to start to disentangle the variables, we would have to randomly assign people to live in large housing estates or lush penthouses. There might be opposition to such a trial. In any case, as scientists we would not be able to make a prediction for a particular child. We may, of course, use population statistics, as insurance companies do, in order to quote the *probability* of a child becoming a future delinquent, given certain genetic and

[6] Lady Capulet said that she was going to contact a hit-man in Mantua to deal with Romeo – but that message didn't get through in time either.

[7] The y-chromosome or testosterone might, however, be arguable.

environmental conditions. This is not the same as understanding the *causes* of delinquency.[8]

Is it possible that understanding the causes of delinquency can be achieved by careful longitudinal studies of many specially selected families? The issues are complex, and I am not an expert in the technical job of examining longitudinal data. I am wary of this approach, however, because longitudinal data is only data about a selection of behaviours, and in an area such as delinquency, different behaviours are found at different ages and in different contexts. Tracing something like an underlying propensity to violence would need many preparatory studies just to define and validate suitable measures. From time to time I will review the work on psychosocial pathways, which seems to lead to an interpretation of *disadvantage* that is different from *cause*.

The need for a framework for thinking in ■

The analysis in the previous two sections has given some hints as to where we are going. The cause of an individual suicide may be impossible to establish definitively, even though we might be able to say more about the contributions of genes and environment to patterns of individual differences in the risk of suicide more generally. But, as the brief analysis of delinquency indicated, we must be careful to distinguish contingency from cause. We have also seen that broad claims about genes or about the role of particular aspects of the environment – the kinds of account beloved of politicians – require more circumspection. Of course, science does not always provide the circumspection required. Newspapers may be quick to produce headlines assigning responsibility for promiscuity, homosexuality and so on to particular genes, but behind a lot of these stories is a scientist who has made a similar claim on the basis of inadequate evidence, and with the underlying message that no other account need be given.

To protect us from error, we need a scientific framework which is suited to the task that we have set ourselves. I explain what I mean by the term 'framework', and how it differs from 'model' and 'theory', in box 1.1.

[8] Of course, such an actuarial exercise may be the appropriate thing to do in some contexts. For example, it would be the appropriate way of estimating the future needs for social services in a particular region. The question of *cause*, in any of its senses, would not be relevant here unless there were a massive break with tradition, with a multi-agency, long-term attack on the problem.

Box 1.1 A note on models and related things

In my thinking, I make distinctions between *models, theories* and *frameworks*. Other people may use these terms differently, so, to avoid unnecessary confusion, I will introduce my own distinctions. The reason why I am stressing this is that you can use the causal modelling *framework* that I am proposing without believing our particular *theory* about, say, autism or dyslexia. Indeed, it seems to be the case that my colleagues and I become more convinced of the potential of alternative causal theories after we have expressed them in a causal modelling notation.

Framework

A framework is a set of ground rules that a community agree on to enable them to express and discuss ideas in a commonly understood fashion. This agreement is usually tacit. These rules would include the types of data that are allowed to influence or test a theory. When people who are operating within different frameworks disagree, it is often because they do not understand each other or because they have conflicting priorities, not because they disagree about the facts or their interpretation. In fact, both people could be correct.* Note that a particular framework may not allow the expression of certain kinds of data. Thus, the framework within which most linguistic theory is expressed does not allow discussion of the time course of speech. This is because linguistic theory is concerned with the underlying structure of language, not with language behaviour. Such facts put a systematic restriction on the range of an individual framework, and, it might also turn out, put restrictions on the scope of the causal modelling framework.

Theory

A theory is an expression of a hypothesized relation among data. Since there is always a choice of data to include in theories, they will always be systematically incomplete. Theories must always be more general in their form of expression than the data that they attempt to encompass. Hence, a theory will always make a prediction about new data. Note that a particular theory can usually be expressed in different frameworks. Thus, two superficially different theories may cover the same data set and make exactly the same predictions about new data. An analogy for this, which might be useful, is the alternative expressions of a circle as $x^2 + y^2 = c^2$ and $r = c$ (for all Ø). Such alternatives can also

be known as *notational variants*. One worked out example of this is the relation among associative nets, schema frameworks and the Headed Records framework in accounting for memory phenomena such as context-sensitive recall (Morton et al. 1985; Morton & Bekerian 1986).

Model

A **model** is a way of presenting a theory. The modelling *method* (such as the one I am going to present) is ideally free of both framework and theory assumptions. The model is often a means of generating predictions from the underlying theory. In practice, the modelling itself reveals assumptions that have been made, so that models often look different from the originating theory. From my own work, an example of this is in the way in which the **logogen model** (Morton 1969) handles the interaction between stimulus and contextual information in the recognition of words in context. The underlying theory merely specified that this interaction took place. The mathematical model required a specification of the nature of the interaction (in fact, the addition of activation).

It is important not to confuse the form of a model with its content. For example, the form of expression of much information processing theory has been that of labelled boxes joined by arrows. The underlying theories are sometimes dismissed as 'mere boxes and arrows models', as though the boxes or the arrows themselves had inherent content that could be true or false. I use boxes and arrows throughout this book, but usually they have a different meaning from that in an information processing model or an information flow model (as in figure 10.3). Of course, anything expressed by boxes and arrows can also be expressed in other ways – including, if you have that particular pathology, words.

*Note that I am not saying that any story is as good as any other story. I happen to believe that one particular story is the correct one. However, that story could be told in a number of different ways, and it would still be the same story. What is told is the science. Where things go wrong is when the limitations of a framework go unnoticed. For example, interpreting the concept of intelligence within a social framework is fine and might help in the formulation of egalitarian policies. However, the inability to represent biologically given individual differences in intelligence within such a framework should not – but, regrettably, sometimes does – lead one to deny the possibility of such individual differences.

What we need is a framework within which *causal* theories can be expressed. What properties would be good to have in this framework? Here are a few:

1 The framework would allow us to represent complex claims about the cause of a disorder in an easily understood manner.
2 It would require us to explicitly distinguish causal relationships from merely contingent ones.
3 It would allow both genetic and environmental factors to enter into any claims.
4 It would distinguish clearly between cognitive and behavioural factors (I will say more about this in the next chapter).
5 It would enable us to represent alternative theories in an easily comparable form – essentially, the framework would be theoretically neutral.

Unfortunately, there does not seem to be such a framework. The nearest are the kinds of representations that have been developed for behavioural genetics and psychosocial pathways. These are both ways of representing the outcome of certain kinds of statistical analyses on population data, and so do not enable us to make explicit distinctions between association (contingency) and cause. So, Uta Frith and I set out to develop a new framework that would help us to think about cause in developmental disorders. We have called this the **causal modelling** framework (Morton & Frith 1995). It is important to note that a framework, of itself, makes no empirical claims about any pathological condition, nor does it commit the user to any particular theory about anything (see box 1.1). This feature makes a framework a neutral forum for the comparison of alternative, or even contradictory, theories. Our aim is that any coherent theory about developmental psychopathology should be expressible within the framework. From this point of view, it doesn't matter if you consider a particular theory to be wrong, or incomplete. If it is coherent, and it is supposed to do with cause, then it will still be expressible within the framework. For example, suppose that there are two claims about a particular condition, one that it has a single genetic cause and the other that there are multiple genetic causes. The consequences of these two alternative claims can be mapped out over biological, cognitive and behavioural levels in ways that enable them to be compared. When the two competing theories are represented in a directly comparable fashion in this way, empirical

data can then be brought to bear on those points that emerge as critical in deciding between alternative theories. If a theory is unclear, however, or inconsistent or simply confused, then our aim is that the framework can be used to isolate the problem, and help to discover possible coherent theories behind the confusion. Put more bluntly, you can use the causal modelling framework to talk about any developmental disorder, even if you disagree with everything that my colleagues have written about that disorder. Indeed, if you want to persuade us that we are wrong in our views about autism or dyslexia, for example, you are more likely to do so within the framework than outside it, since, inside the framework, you can be more certain that we will understand your ideas and accept that your causal arguments are coherent. But let me stress again that my aim in this book is not to debate theories but to examine ways of representing (modelling) theories.

Creating a tool: the problem of notation

In the course of establishing diagnostic categories, and in the course of attempts to explain developmental psychopathology, a variety of claims are made that touch on the principles I have mentioned above. The debate between proponents of opposing views is often confused. There is a lack of clarity of expression, as well as much unresolved conflict. There is sometimes even more conflict between people whose ideas turn out to be minor variations on each other. There is a sociological account of some of this – to do with the fury of competition for the same piece of turf – but much of it, I believe, is because the only tool that most people have for communicating their ideas is language. Purely verbal expression of immensely complex ideas is difficult to achieve. I find that such expression of ideas is even more difficult to comprehend. This is because language is predominantly linear, while ideas are multidimensional in their relationships. I propose that some of the problems of understanding what is going on can be relieved by use of a graphical notation within which the underlying ideas can be expressed. A graphical representation of ideas can reveal structure that was previously obscured.

This is really a very simple idea. Take the following problem:

> Jim was sitting on Helga's left. Helga was opposite Mary, who was between Pierre and Lorraine. Pierre was on Jim's left. Who was opposite Lorraine?

Such a problem is difficult to solve without constructing a diagram of the seating plan:

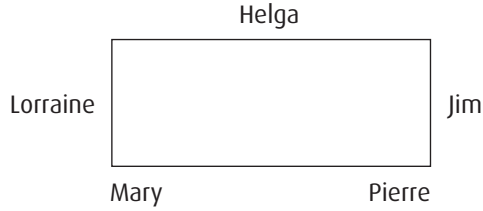

With such a plan, much more difficult problems become child's play.

An example of the limits of language – be careful when you read this ■

Let me begin to illustrate the limitations of language with an example from our own work. In Morton and Frith (1993a), we comment on some of the implications of a paper by Cossu et al. (1993). These authors showed that Italian children with Down's syndrome could read non-words and yet failed on certain phoneme segmentation tasks. These are tasks in which the children are asked to play games with sounds. For example, they might be asked to delete the first sound in the word 'table' – in which case the correct answer is 'able'. These two facts, it was claimed, contradicted certain theories about reading acquisition. In commenting on this paper, Uta and I found that we needed to specify some of the cognitive abilities that were necessary if the children were going to carry out particular tasks successfully. To start with, what ability was necessary for the child to be able to succeed on typical phoneme segmentation tasks? Let's say that successful performance on such tasks requires both the development of a **grapheme–phoneme** (**GP**) correspondence system and a competence in relation to phonemes, which could be called **implicit phoneme awareness** (**iPA**). The GP system would also be a prerequisite for non-word reading tasks and the iPA system would be needed for the understanding of rhyme. Both GP and iPA depend on a common underlying phonological system, P. The factor M, which underlies meta-representational skills, is a prerequisite for iPA (but not GP). We continue:

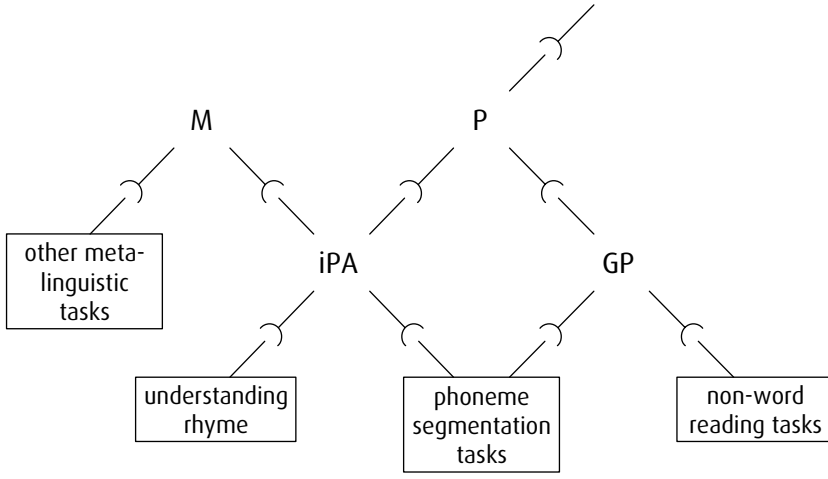

Figure 1.6 A developmental contingency model (Morton 1986) to account for the data in Cossu et al. (1993). The horseshoe shape between two elements is to be understood in the following way: the former is necessary for the normal development of the latter. The detail of the diagram might be intelligible in conjunction with the text (from Morton & Frith 1993a).

> it can be seen that failure in phoneme segmentation tasks would result from lack of development of iPA. This could be due to an absence of P (which would also cause a lack of GP) or from absence of M. We assume that this is the case for the Down Syndrome children and predict that they fail both on iPA dependent tasks and on other metalinguistic tasks. (Morton & Frith 1993a, p. 295)

I predict that most of you will not have too easy a time with the above passage. In fact, in the journal article, the prediction was accompanied by five diagrams, the last of which is reproduced in figure 1.6. The diagrams make the descriptions seem self-evident, even if you are unclear about the nature of the components.

For a second example, take the following passage from Russell (1996). This is a little unfair, since the passage in question was preceded by 250 pages of closely argued text concerning the nature of the cognitive deficit in autism.

> If agency plays the role in the development of self awareness – or 'ego development' – that I have claimed it does, then human beings with congenital impairments in agency will undergo deviant ego-development;

> and if ego-development is impaired so too will be the acquisition of a theory of mind. My concern here is with what I called 'mind$_o$' – the individual's immediate knowledge of being a centre of responsibility for his or her action-generated experiences, of having the power to determine what is experienced and thus of being the 'owner' of a mental life. If there is an early impairment in the ability to determine what one does and how one's attention is directed then there will develop an inadequate sense of oneself . . . (Russell 1996, p. 253)

I think that this causal model would be considerably easier to appreciate with a diagram. One of the problems is that Russell uses different words to refer to the same conceptual entity. Thus, we have

> the individual's immediate knowledge of being a centre of responsibility for his or her action-generated experiences

meaning the same as (having the same reference in the argument as)

> the ability to determine what one does

and

> the power to determine what is experienced

referring to the same idea as

> the ability to determine how one's attention is directed

This is an interesting literary convention, not to fatigue the reader with repetition of word or phrase, but for certain kinds of modelling and deductive thinking one needs to be sure that the referent has not changed in the smallest way. We are all familiar with the kind of creeping referent that can occur in political settings, where the argument starts with, say, children with mothers who work, moves to children who play in the street, thence to juvenile delinquents and finally to psychopathic murderers.

Words can be sharp and precise: they convey nuances in a way that a newly invented visual notation cannot possibly offer. I wonder, however, whether in scientific debates the sharpness of verbal definitions sometimes exceeds the refinement of the underlying thought. For this reason, we often hear the cry that somebody's new theory is 'just words'. I propose that a notation that has not yet gathered bad habits, sloppy

practice or ossified schemas is useful to check out the soundness of verbally formulated theories. In no way do I claim that a theory can or should *only* be visually represented, but we have to beware of purely rhetorical science.

There are other reasons too for opting for a visual instead of verbal notation. Words have a short shelf life. They change with use, and many a discussion ends with the realization that the discussants have used the same words with entirely different meanings. The use of a suitable framework will enable us to anchor such terms in relation to particular kinds of behaviour, tasks or experiments.

An invitation to consider diagrams as a tool

We have come to believe that there are wide individual differences in the way in which people use words or diagrams in scientific texts. Doubtless, there are a number of reasons for this. Some of the individual differences might be related to individual differences in people's abilities to use maps, or differences in skill with verbal reasoning. Certainly, I have difficulty in fully comprehending any complex scientific argument unless I have pencil and paper in hand to help convert it into diagrammatic form. In addition, however, there will be factors such as familiarity or simple prejudice. Clearly, it also depends on training. It's hard to imagine mathematics, physics, anatomy, physiology or biochemistry without diagrams, although one could use a thousand words instead for each ('the hip bone's connected to the thigh bone . . .'). It is certainly the case that behavioural scientists are used to mere words, and many a textbook on psychology can be found without a single diagram, apart from graphs. But I regard this as a pathology in itself, as a symptom of something even worse – such as behaviourism or reductionism.[9]

A tool for representing causal relationships

Cause is a complex issue in any domain. An additional complication is that, as I will stress, the set of causal elements (the terms used in tracing causal chains) come from different types of thing – from genes, through brains and cognitive entities to socio-economic factors. So we will need to

[9] The reasons for this leap are complex and based on experience. I leave it to historians and philosophers of science to spell out the inferential links.

be able to express competing claims that some problem has been inherited, caused by brain damage at birth, by pollution in the environment, by lack of learning, by bad parenting or by peer pressure, or through the interaction of genes with the environment. We may also want to refer to different time scales. Some developmental disorders have their roots in the genes, and we may want to refer to elements in the causal chain whose existence is brief, because they play only a transitional but crucial role in neural development. On the other hand, we will want to refer to the immediate cause of some current behaviour as based on a combination of some aspect of the environment in interaction with an abnormal or deficient process that is specified at the cognitive level.

What I am presenting in this book is a method of helping to express causal statements that include any and every kind of element. In the method itself, there is nothing to believe or disbelieve – it is simply a tool. I believe that this tool, which we call *causal modelling*, is more than useful. I believe that it can help to raise standards in evaluating different ways of looking at etiology, and in diagnosis. However, it will not change the basis of your thinking. As a tool, it has the properties of a tool. You do not live by tools. You pick them up and use them when you need to. Physical tools help you to do things that you could not do by yourself. Conceptual tools serve exactly the same function, but because they are conceptual they help you mentally – remind you; suggest ways to do things; start you off; help you to know when the job is finished or characterize what is left still to do; and tell you what other tools or equipment (such as facts or experiments) are needed. There are a number of features to watch out for in the causal modelling framework:

- it is a tool that helps you to clarify your ideas
- it is a memory tool – it enables you to represent all the complexity you want to represent in an easy-to-understand form
- some checks for inconsistency are built in
- it enables you to establish both common ground and incompatibility with others with a degree of precision
- it allows the formulation of more precise tests of alternative ideas
- it reveals what you could know but don't yet know.[10]

I also hope that the book will help you think about development and developmental disorders in a new and more productive way.

[10] For instance, the gap between a chromosomal abnormality such as trisomy 21 and the anatomical condition of mongoloid eye folds in Down's syndrome.