Understanding Developmental Disorders

A Causal Modelling Approach

John Morton

Understanding Developmental Disorders

For G and Pooch, with love

Understanding Developmental Disorders

A Causal Modelling Approach

John Morton

BLACKWELL PUBLISHING
350 Main Street, Malden, MA 02148-5020, USA
9600 Garsington Road, Oxford OX4 2DQ, UK
550 Swanston Street, Carlton, Victoria 3053, Australia

First published 2004 by Blackwell Publishing Ltd

2 2006

Library of Congress Cataloging-in-Publication Data

Morton, John, 1933–
Understanding developmental disorders : a causal modelling approach / John Morton.
p. ; cm.
Includes bibliographical references and index.
ISBN 0–631–18757–X (hardback : alk. paper) — ISBN 0–631–18758–8 (pbk. : alk. paper)
1. Developmental disabilities—Etiology. 2. Child development deviations—Etiology. 3. Conduct disorders in children—Etiology. 4. Cognition disorders in children—Etiology. 5. Mental retardation—Etiology. 6. Parent and child. 7. Cognition. 8. Causation.
[DNLM: 1. Developmental Disabilities—physiopathology—Child. 2. Developmental Disabilities—psychology—Child. 3. Models, Theoretical. WS 350.6 M889 2005] I. Title.

RJ506.D47M67 2004
616.85′88071—dc22

2004009013

ISBN-13: 978-0-631-18757-8 (hardback : alk. paper)
ISBN-13: 978-0-631-18758-5 (paperback : alk. paper)

A catalogue record for this title is available from the British Library.

Set in 10/12.5pt Baskerville
by Graphicraft Limited, Hong Kong
Printed and bound in Singapore
by C.O.S. Printers Pte Ltd

The publisher's policy is to use permanent paper from mills that operate a sustainable forestry policy, and which has been manufactured from pulp processed using acid-free and elementary chlorine-free practices. Furthermore, the publisher ensures that the text paper and cover board used have met acceptable environmental accreditation standards.

For further information on
Blackwell Publishing, visit our website:
www.blackwellpublishing.com

Contents

Preface and Acknowledgements viii

Chapter 1 Introducing Cause 1
Cause and public issues 1
Cause and individual events: 'Why did Romeo die?' 6
Some more reasons for not looking at individual cases 9
The need for a framework for thinking in 10
Creating a tool: the problem of notation 14
An example of the limits of language 15
An invitation to consider diagrams as a tool 18
A tool for representing causal relationships 18

Chapter 2 Introducing Cognition 20
One thing I do want you to believe 20
Reductionism 22
Can we rely on behaviour? 24
The IQ example: a note of caution 27
Why cause needs cognition 29

Chapter 3 Representing Causal Relationships: Technical and Formal Considerations 34
Categorizing facts 34
The causal notation 38
Starting a causal model for autism 41
Complications 46
Some easy stuff on cause and correlation 51
Other notations 54

Chapter 4 Autism: How Causal Modelling Started **67**
The biological origin of autism 74
The role of cognition in defining autism 81
What is mentalizing? 86
The non-social features of autism: how to diagram ideas on weak central coherence in autism 89
Summary 92

Chapter 5 The What and the How **98**
Ground rules of causal modelling 99

Chapter 6 Competing Causal Accounts of Autism **106**
Representing the effects of environmental factors 107
Cognitive theories of autism 112

Chapter 7 The Problem of Diagnosis **133**
Diagnosis and cause: relying on behaviour 134
The Spanish Inquisition example: the dangers of labelling 135
Problems of diagnostic practice 140
Variability 148
Changes over time: improvement and deterioration 152
The variability of the phenotype 153
On co-morbidity and the question of residual normality 158
To summarize 160

Chapter 8 A Causal Analysis of Dyslexia **161**
The dyslexia debate: Is there such a thing as dyslexia? 161
The discrepancy definition of specific reading disability 164
Towards a cognitive definition 166
An X-type causal model of dyslexia 168
Competing theories of dyslexia 176
Non-biological causes 195
Other biological causes of reading failure 199
How do we sort among the options? 200
The relationship between acquired and developmental dyslexia 204
A theoretical update 204

Chapter 9 The Hyperkinetic Confusions **208**
Drugs as diagnostic refinement 212
Types of theory 216
The problem of co-morbidity: conduct disorder and ADHD 218
The cognitive level 219
Sonuga-Barke's dual pathway model 223
Summary 226

Chapter 10 Theories of Conduct Disorder **227**
The violence inhibition mechanism (VIM) model 228
The social information processing model for aggressive children 231
The coercive parenting model of Patterson 235
The theory of life-course persistent antisocial behaviour 236
What does the application of the framework tell us about the theories? 244

Chapter 11 Tying in Biology **247**
Relations between the cognitive and biological levels 247
Equivalence: brain to cognition 251
Causal influences from cognition to brain 253
Genes and cause: the end of behaviour genetics 255
Endophenotypes 264
Mouse (and other) models for human disorders 266

Chapter 12 To Conclude **270**

References 273
Name Index 292
Subject Index 296

Preface and Acknowledgements

Human development can be seen as the unfolding of a particular pattern. Genes are responsible for producing proteins that lead to structures of various kinds – legs and hands and hearts and tongues and, the bit that concerns us, brains. Within the brain there is structure, and this is formed under the influence of the genes in complex ways, together with the influence of the environment, both internal and external. There is a lot of variation in the genes, and a lot of variation in the environment, but the result is a human being. The variations in the genes and the variations in the environment give rise to individual differences – in eye colour, height, temperament and intelligence – within the range of what is often termed normality. This is a difficult term, however, since what is normal is often confused with what is acceptable. A man who is seven foot tall is acceptable in the context of a professional basketball team, but not in the context of an airline seat, particularly next to me!

Things can go wrong with early brain development: there might an unfortunate combination of allelic variants of genes; one of the critical genes might be damaged; the environment might be unsuitable – for example, the intra-uterine environment provided by a drug-taking mother. As result of one of these misfortunes, the child develops outside the normal range – the problem may be a general one or may be specific, but in either case some health or education professional is going to want to describe the condition and relate that child to other children who are similar. The reason for doing this may be to help with prognosis – to predict what the future may be like for that child. The description will be of use in helping to understand the nature of the condition, predicting other kinds of problems, and giving a hint as to

what might happen in various situations or in the future. It will also be help to guide any treatment plan.

Later on, the interaction of the child with his or her environment may make things better or worse. In some cases, the environment may be so aversive that children with physically normal brains may develop in a way that leads them to be termed abnormal in some way. Such children will also attract labels.

Sometimes, the labels that are used relate only to the child's behaviour – conduct disorder is a good example. Sometimes the labels refer to a specifiable genetic problem – William's syndrome is an example here. In other cases the labels have deeper significance, with the defining features being more complex – autism is a good example here. One of the problems with the labels is that different people mean different things by the same label. Take 'dyslexia', for example. For some people, this is simply a term applied to children whose reading age is significantly below their mental age (usually by two years). For others, it means people with a deficiency in their magnocellular cortex. For yet others, it means people with a particular deficiency in phonological processing (and who, curiously, may have no observable problem in reading!). How can we discuss such different types of theories in the same breath?

I have felt for many years that understanding developmental disorders requires us to think about biological, cognitive, behavioural and environmental factors, and to discover the causal relationships among these elements. This is a way of thinking about developmental disorders which grew from my time at the Medical Research Council's Cognitive Development Unit (CDU) in London. My colleague Uta Frith and I felt the need for clarity in theorizing about this process and developed a methodology for enabling that clarity. We called the method 'causal modelling', and wrote a very long chapter about it (Morton & Frith 1995). This book is an extension of that work. This book is about the nature of causal theories.

The actual theories that I use are partially chosen for illustrative reasons, as I will repeatedly remind you. What I cannot do here is to list all the most up-to-date theories of each of the developmental disorders. To give a complete up-to-date account of any one of autism, dyslexia, conduct disorder or ADHD would take a book of the length of this one. So I have taken the theories I am most familiar with, and those that illustrate the range of theories that exist. I don't think that use of the causal modelling method makes any of these theories more correct.

Instead, it helps to identify weak spots in the theories, and also helps us to see the relationships between different theories. This makes it easier to set up ways of testing the theories against each other. This is what the book is about. At the CDU, we did develop our own beliefs about autism and dyslexia, for example, and, although I have tried to be dispassionate I may seem to favour those beliefs here. However, these beliefs are not fixed, and it is important for the reader to realize that what I want to communicate in this book is a way of thinking about different theories of developmental disorder. I am proposing a tool. This tool has been found useful by students and by practitioners, as well as by fellow scientists. It is offered as a way to help you formulate your own theory of any developmental disorder as well as a way of understanding other people's theories.

Of course, you cannot just pick up a complex tool and use it; you need to think about the problem appropriately. Using causal modelling encourages you to think about development in a particular way. Specifically, you have to learn to work productively with cognitive concepts, and clarify in your own mind the distinction between cognition and behaviour. It was the application of these principles to the problems of developmental disorders that I learned from Uta Frith. In a sense, causal modelling was my attempt to formalize the way in which Uta thought about autism and dyslexia. Her voice will be heard often through the book, and the only reason she is not co-authoring it is that she is too busy moving on our understanding of these disorders. I hope that she endorses much of what I have written, but I excuse her from endorsing it all.

The book has chapters of two kinds, those concerned with setting up and explaining what causal modelling is and how to do it, and those chapters illustrating how causal modelling can help us to understand and compare different theories of particular developmental disorders. Since this is a book on causal modelling, I will rarely refer to any theory that does not make causal claims about the syndrome. The major exception is in the chapter on conduct disorder (chapter 9), where, following Krol et al. (2004), I look at the structure of the main theories of conduct disorder, which are not causal in nature.

It is worth noting briefly the contrast between causal modelling as I have been describing it here and structural equation modelling, which is also sometimes called causal modelling (for an overview, see Fife-Shaw 2000). Structural equation modelling is an analytical technique that is applied to data. Roughly speaking, you look for *latent variables*

that mediate between measured variables. There may be an *a priori* theory connecting the variables, but the techniques do not depend on this. All you need know is that this is not what I am going to talk about.

In chapter 1, I introduce the notion of cause, comparing the everyday use of the term with a more scientific use. This discussion leads us to appreciate that when we are trying to account for a developmental disorder we have to distinguish between internal and external factors and, within the internal factors, we have to distinguish between biology, cognition and behaviour. I then discuss the need for a scientific tool, a tool to aid thinking, a tool which uses diagrams instead of words. Here I point out for the first time, as I will point out repeatedly, that use of the tool does not commit one to any particular theoretical position.

In chapter 2, I spend some time explaining why the study of developmental disorders requires us to consider cognition. This is the only theoretical point that I argue for in the whole book. If you don't believe that cognition has any role to play in the study of developmental disorders, and if I fail to convince you in this chapter, then you had better stop reading and I will give you your money back. But your reasoning had better be good!

The elements of causal modelling are introduced in chapter 3. The fundamental principle is that the verbal expression of your theory gets converted into a diagram. The notation is very simple, and I work through three examples of different types to give you a feel for what I am trying to do.

Chapter 4 is of the second kind. In it I take you through some of the history of the development of causal modelling – the evolution of the 'mentalizing deficit' account of autism.

In chapter 4, a number of issues raise themselves about the causal modelling technique. In chapter 5 I bring these issues together in the form of a number of maxims. These maxims go towards helping to establish good practice in formulating theories of developmental disorder and representing them in a causal model.

In chapter 6, I go through most of the current alternative theories of the origin of autism and frame them using the principles of causal modelling. This exercise helps us to see the relationships among the various theories, and pinpoints which theories are insufficiently explicit or are ambiguous in expression.

One of the major issues which becomes evident in chapter 4 is that of diagnosis. Thus, if we do not know what autism is we cannot have any theory about its development, nor can we properly diagnose it.

In chapter 6, we face these issues and show how causal modelling can help us to reconcile different approaches to diagnosis.

Chapter 7 takes up some of the general points raised in the previous chapter, concerning the notation of causal modelling and the ways that are available for elaborating our causal theories.

In chapter 8, I apply the lessons we have just learned to a discussion of the various current theories of dyslexia. In particular, the causal modelling method enables us to show clearly how some theories can be embedded in other theories, rather than being directly opposed. In this way, the debate about dyslexia can be focused on specific questions rather than involving major confrontations.

In chapter 9, I look at theories of ADHD. More properly, this is a set of disorders that are usually defined in terms of behaviour, without regard to the origins of the behaviour. This has led to much confusion over the years and it is only recently that decent cognitive theories have emerged.

Chapter 10 is devoted to a detailed analysis of conduct disorder. This work, based on a collaboration with Nicole Krol and Eric de Bruyn of Nijmegen, focuses on a disorder to which there is thought to be a major contribution of environmental and developmental factors. In this respect, the disorder is different from autism and dyslexia, for which there is a dominant well established biological/cognitive base. In this chapter I also compare causal modelling with other ways of talking about developmental disorders, particularly psychosocial pathways.

In chapter 11, I look at some of the more general issues of the relationships among the three levels of description; in particular, the relation between brain and cognition. The use of causal modelling has sharpened my own view of this relation.

Now let me say two words concerning what this book is not about. To start with, it is not about the nature of cognitive mechanisms. I do not address the question: How does cognition cause behaviour? This is not my concern here. I will say more about this in chapter 2. Nor will I say much on the topic of the relationship between a developmental disorder and an acquired disorder, although I will touch on that topic in chapter 8.

Developmental psychology is difficult, more difficult, I think, than constructing theories of adult cognitive function. I never managed to become a competent developmentalist, but managed in the CDU environment by being able to bring theoretical methods from the world of cognitive psychology, and by having support and continual interaction

from a number of outstanding developmental scientists. Their influence pervades this book. I have already mentioned Uta Frith, and I should add Annette Karmiloff-Smith, the late Rick Cromer, Alan Leslie, Mike Anderson and Mark Johnson. Of the many friends, visitors and short-term staff that we had at the CDU, Sue Carey, Anne Christophe, Mike Cole, Francesca Happé, Lila Gleitman, Mani das Gupta, Geoff Hall, Kang Lee, Jean Mandler, David Olson, P. Prakash and Prentice Starkey all played a role in my education.

In relation to the manuscript of the book, thanks are particularly due to Sarah-Jayne Blakemore, Elisabeth Hill, Franck Ramus and Essi Viding, who I have encouraged to be remorseless and detailed in criticism and who have repaid me handsomely! I have not always taken their advice, and no doubt will regret it. Mike Anderson, James Blair, Nichole Krol and Edmund Sonuga-Barke have also made invaluable suggestions. Sarah White has helped in many ways, particularly organizational.

Most of all I am indebted to my wife, Guinevere Tufnell, who has not only tolerated the five-year incubation of the book but has added invaluable clinical insight throughout.

The figures were drawn using CHARTIST™.

John Morton

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?

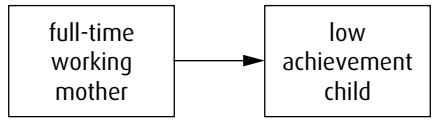

Figure 1.1

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

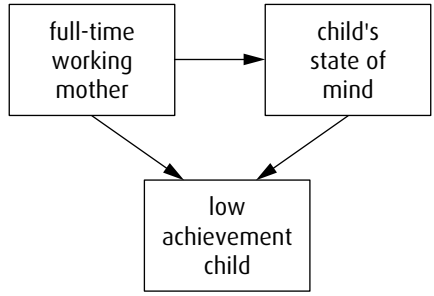

Figure 1.2

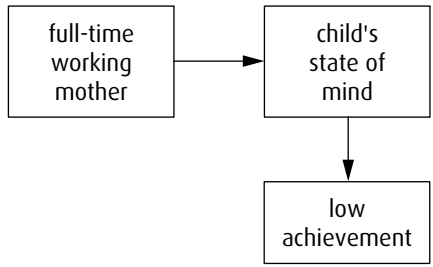

Figure 1.3

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

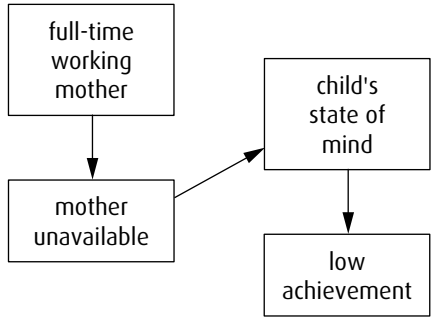

Figure 1.4

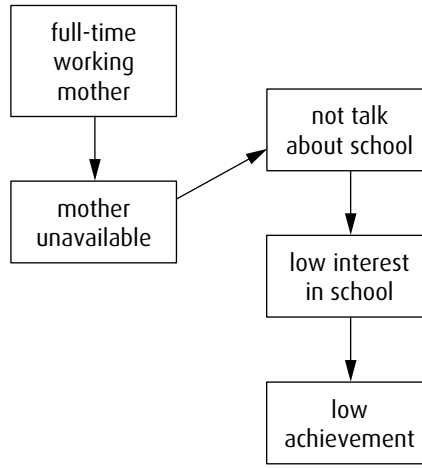

Figure 1.5

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:

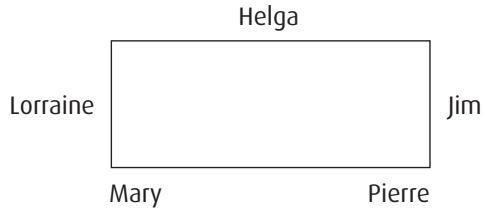

With such a plan, much more difficult problems become child's play.

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: