

1

Understanding Developmental Cognitive Disorders

John, Peter, and Ann are three 7-year-old children. John's parents and teachers have concerns about his progress in learning to read. John is generally bright and understands concepts well. Formal testing showed that he had a high IQ (120) with somewhat higher scores on the performance than the verbal scales of the test. John could only read a few simple words on a single word-reading test – a level of performance equivalent to a typical 5½-year-old child. John does not know the names or sounds of several letters of the alphabet. Verbally John is a good communicator, though he does show occasional word-finding problems and occasionally mispronounces long words. John is a child with dyslexia.

Peter is also a bright little boy (IQ 110, but with markedly lower scores on the performance than the verbal subtests). He has made a very good start with learning to read, and on the same test given to Peter he read as many words correctly as an average 8-year-old child. Peter has severe problems with games and sport at school, particularly with ball games. He is notably ill-coordinated and frequently drops and spills things. He has very serious difficulties with drawing and copying, and his handwriting is poorly formed and difficult to read. Peter has developmental discoordination disorder.

Ann is a socially withdrawn child. She avoids interacting with other children in school whenever she can. She is sometimes observed rocking repetitively and staring out of the classroom window. Ann's communication skills are very poor, and she appears to have quite marked difficulties understanding what is said to her, particularly if what is said is at all abstract. When an attempt was made to give Ann a formal IQ test, testing was discontinued because she refused to cooperate. The few items she did complete suggested she would obtain a very low IQ score. Ann is fascinated by cars, and will spend many hours cutting out pictures of them to add to her collection. Ann is a child with autism.

These three cases of 7-year-old children illustrate some of the varied cognitive problems that can be observed in children. In this book we will attempt to provide a broad survey of the major forms of cognitive disorder found in children, and lay out a theoretical framework for how these disorders can best be understood.

2 Understanding Developmental Cognitive Disorders

Understanding these disorders, in turn, holds prospects for how best to treat them. Our approach to these disorders is from a developmental perspective, by which we mean that a satisfactory understanding of these disorders needs to be informed by knowledge of how these skills typically develop. Most of the explanations we consider in the book will focus on the cognitive level: a functional level dealing with how the brain typically learns and performs the skills in question. Wherever possible, however, we will relate these cognitive explanations to what is known about the biological (genetic and neural) mechanisms involved in development. The interplay between genetic, neural, and cognitive explanations for behavioral development is currently an area of intense activity and excitement.

Some Terminology for Classifying Cognitive Disorder

In this book we will consider a wide range of developmental disorders that affect language, learning, and cognition. The disorders considered include those affecting language, reading, arithmetic, motor skills, attention, and social interaction (autism spectrum disorders). There are a number of features that are shared by the disorders we will discuss: they all occur quite commonly and have serious consequences for education, and thereafter for well-being in adulthood; and there is good evidence that all these disorders reflect the effects of genetic and environmental influences on the developing brain and mind.

To begin with it is important to distinguish between specific (or restricted) difficulties and general difficulties. Specific difficulties involve disorders where there is a deficit in just one or a small number of skills, with typical functioning in other areas. General difficulties involve impairments in most, if not all, cognitive functions. Terminology in this field differs between the UK and the USA; we will consider both here, but we will use primarily British terminology in later sections of the book.

In the UK a selective difficulty in acquiring a skill is referred to as a “specific learning difficulty.” The term learning difficulty makes it clear that skills must be learned; specific means that the difficulty occurs in a restricted domain. Dyslexia is one of the best known and best understood examples of a specific learning difficulty. Children with dyslexia have specific difficulties in learning to read and to spell, but they have no particular difficulty in understanding concepts and may have talents in many other areas such as science, sport, and art. In the USA (following DSM-IV, the *Diagnostic and Statistical Manual of Mental Disorders* of the American Psychiatric Association) such specific difficulties are called learning disorders.

Specific learning difficulties can be contrasted with general learning difficulties (or, in US terminology, mental retardation). General learning difficulties involve difficulties in acquiring a wide range of skills. People with the chromosomal abnormality of Down syndrome, for example, usually have general learning difficulties and typically have problems in mastering all academic skills and with understanding in most domains. In this book we will focus upon specific learning difficulties.

In practice, the distinction between specific and general learning difficulties is often based on the results of a standardized IQ test. IQ tests (or measures of general intelligence) are highly predictive of variations in attainment in all manner of settings. The average IQ for the population is 100 (with a standard deviation of 15 points). In the UK people with IQ scores between 50 and 70 are referred to as having moderate learning difficulties, and people with IQ scores below 50 are said to have severe learning difficulties. US terminology distinguishes between mild (50–70), moderate (40–50), severe (25–40), and profound (IQ below 20) mental retardation. Often the diagnosis of a specific learning difficulty is made only in cases where the child achieves an IQ score in the average range (perhaps an IQ of 85 or above).

Operationally the distinction between specific and general learning difficulties is therefore quite clear: children with specific learning difficulties typically have average or near to average IQ scores, while children with general learning difficulties have IQ scores below 70. Conceptually, however, the distinction is probably a bit more slippery. It is important to appreciate that there is a continuum running from the highly restricted deficits found in some children (e.g., a child with a severe but isolated problem with arithmetic), to more general difficulties (e.g., a child with severe language difficulties who has difficulties both with understanding speech and expressing himself in speech), and to very general difficulties (a child with an IQ of 40, who is likely to have problems in reading and spelling, as well as spoken language, together with a range of other problems including problems of perception, motor control, and general conceptual understanding). One aim of this book is to convey an appreciation of how studies of children with different types of learning difficulties have contributed to an understanding of how a range of different brain systems are involved in learning. The range of learning difficulties that occurs ultimately helps us to understand how the developing mind is organized and how the skills that are impaired in some children are typically acquired.

Levels of Explanation in Studies of Developmental Cognitive Disorders

What form of explanation can we hope to achieve for developmental cognitive disorders? It is important to distinguish between the different levels of explanation that are possible. Morton and Frith (1995) have laid out very clearly the logic and importance of distinguishing the different levels of explanation that are needed for understanding developmental disorders. They show how it is essential to consider three major levels of explanation: biological, cognitive, and behavioral. At each of these levels underlying processes (in the child) interact with a range of environmental influences to determine the observed outcome.

We can illustrate the role of different levels of explanation with reference to conduct disorder, a disorder of socio-emotional development that we will not deal with further in this book. Conduct disorder is a disorder where there have been advances in understanding at several different levels recently (Viding & Frith, 2006) and it is

4 *Understanding Developmental Cognitive Disorders*

therefore a good example to illustrate the different levels of explanation involved in the study of developmental disorders. Conduct disorder is defined in DSM-IV as persistent antisocial behavior that deviates from age-appropriate social norms and violates the basic rights of others (American Psychiatric Association, 1994); alternative terms sometimes used for this disorder include antisocial behavior and conduct problems. A model for one aspect of conduct disorder – reactive aggression – proposed by Viding and Frith (2006) is shown below.

This model represents processes operating at the biological, cognitive, and behavioral levels of explanation. It appears that at the biological level specific differences in genes that regulate the action of the neurotransmitter serotonin are important in giving rise to a predisposition to commit acts of violence. More specifically, different variants (alleles) of a gene coding for monoamine oxidase inhibitor A (MAOA) have been indentified, with either high (MAOA-H) or low activity (MAOA-L). Research has suggested that having the MAOA-L gene may predispose an individual to display violent behavior but only if they experience maltreatment in childhood (Caspi et al., 2002). (This is a very important finding since it provides an example of gene–environment interaction; neither having the gene nor being maltreated alone may be sufficient but both factors together give a greatly increased risk of developing conduct disorder.) These genetic and environmental risk factors in turn appear to operate on the development of brain systems concerned with the regulation of emotion. In particular it is thought that the MAOA-L gene may be associated with the development of hyperresponsivity of the amygdala during emotional arousal coupled with diminished responsivity of areas of the prefrontal cortex that normally play a role in regulating such emotional responses. This pattern of brain dysfunction might be seen as providing the biological basis for reacting excessively emotionally and violently when provoked by certain environmental conditions (in everyday terminology, losing control or “losing it” when provoked).

Viding and Frith suggest that these brain differences express themselves at the cognitive level via a mechanism called an emotional intent encoder, which in turn is associated with a bias to fight. Interestingly, in this model, Viding and Frith explicitly propose that the interactive effects of childhood maltreatment operate at a cognitive level by leading to the creation of many emotionally charged memory representations. This is an interesting and testable hypothesis, but of course such effects may also operate at a biological level as well as, or instead of, at the cognitive level.

The final level in the model is the behavioral level, where the fight response bias mechanism may lead to reactive aggression (fighting when provoked) as well as impulsive violence.

A complete explanation of any disorder will involve at least three levels of description. For one aspect of conduct disorder – reactive aggression – genes appear to contribute powerfully to the risk of developing the disorder in interaction with specific environmental experiences (maltreatment) in childhood. It appears that these genetic effects in turn affect the development of brain circuits concerned with the experience and regulation of emotion, perhaps particularly anger, which, in interaction with memories of previous experiences associated with violence, may lead to a bias toward fighting (rather than running away or being afraid). At a behavioral

level, this bias toward a fight response may lead to the observed profile of responding violently when provoked and occasionally committing unprovoked, impulsive acts of violence.

Morton and Frith (1995; Morton 2004) argue that it is useful to make explicit diagrams of these sorts of theoretical explanations, using an approach they term causal modeling. The Viding and Frith diagram (Figure 1.1) is an example. It is important to note that the arrows in such a diagram represent hypothetical causal links. According to this model, a genetic difference causes a brain difference (abnormality), which in turn causes cognitive (emotional) deficits, which in turn cause the observed behavioral patterns (a propensity to violence). Note that within this framework environmental effects can be thought of as operating at each level. So, for example, a virus or early brain injury might also lead to the brain abnormality underlying the emotion control problem, and the effects of positive experiences (a nurturant nonaggressive parental style) might prevent the development of the emotion regulation deficits. Some forms of treatment (teaching anger management strategies) might also have effects on the behavioral level (inhibiting violent outbursts) without having a direct effect on the cognitive level (the person may still feel angry and feel the urge to lash out, but develop ways of controlling such feelings).

It is important to emphasize that all three levels of description are useful, and each helps us to understand the disorder. While links can and should be made between these different levels of explanation, we cannot reduce or replace one level of explanation with a lower level. The cognitive level of explanation (emotion encoding) cannot be replaced by a neural explanation (problems with the amygdala). We would note here that we have followed Morton and Frith's terminology by referring to the level between the brain and behavior as "cognitive." This might seem too narrow a term because cognition essentially refers to thought processes. We will stick with this term for the moment, though in some of the disorders we consider later (as well as in the case of conduct disorder) this terminology might usefully be broadened to consider other forms of mental processes, particularly emotional and motivation processes, that probably cannot simply be reduced to cognition. The point, however, is that we need a level of "mind" or "mental process" as an intervening level of explanation between brain and behavior. We would also argue, in light of recent advances in our understanding of developmental disorders, that the causal model presented in Figure 1.1 is too unidirectional to capture the truly interactive nature of development. It is also necessary to postulate causal arrows running "backwards" from lower levels to upper levels. This at first seems counterintuitive, but some examples help to explain why it is necessary.

Can changes at the behavioral level alter things at the cognitive level? Almost certainly yes. If we take the example of teaching anger management strategies mentioned above, it may be that such training will work by modifying the cognitive mechanisms associated with emotional encoding; seeing a person grin could be interpreted simply as showing that they were happy rather than indicating they are intending to insult you. Do such changes at the cognitive level depend upon changes in underlying brain mechanisms? Again it would seem likely that they do. Connections

6 Understanding Developmental Cognitive Disorders

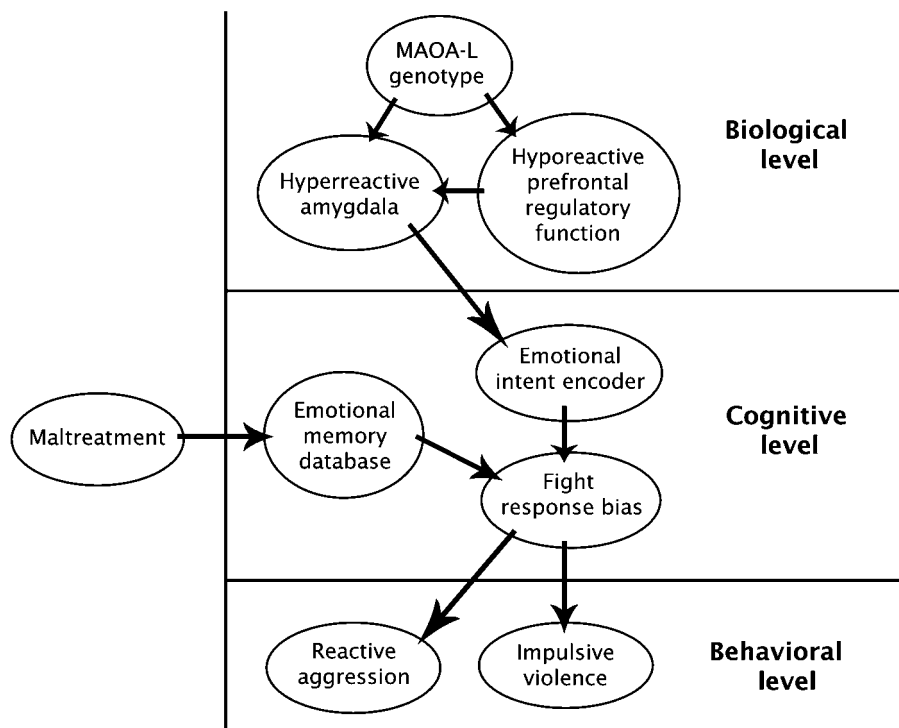


Figure 1.1 A causal model of the potential gene–brain–cognition–behavior pathways from MAOA-L to reactive aggression. (Adapted from Viding, E. & Frith, U. Genes for violence lurk in the brain. Commentary. *Proceedings for the National Academy of Sciences*, 103, 6085–6086. Copyright (2006) National Academy of Sciences, USA.)

between nerve cells may be modified by experience and this in turn will result in lasting structural and functional changes in the circuits responsible for encoding and regulating emotion.

Finally, and perhaps most surprisingly, we can consider whether changes at the behavioral and cognitive levels can affect things at the genetic level. Most people would probably doubt this proposition. Our genetic makeup is fixed (we inherit our DNA at conception and experiences are not going to alter it), this is true, but there is evidence that experiences can alter the way genes are expressed. Genes (genes are sequences of base pairs in DNA) do not regulate development directly. Rather, genes control the production of messenger RNA (mRNA), and mRNA in turn controls the production of proteins in cells. Furthermore, mRNA molecules degrade quickly so that if more of a protein is needed the cells concerned have to keep manufacturing more mRNA. Changes in the rate at which a gene produces mRNA will therefore result in changes in the rate at which the protein coded is produced in a cell. The levels of regulation in cells, as currently conceptualized by molecular biologists, are shown in Figure 1.2. Once again, in this diagram there are different levels of explanation: the genome (the genes that consist of sequences of base pairs in DNA), the transcriptome (the mRNA produced under the control of the base sequences in




Figure 1.2 Diagram showing the complexities of genetic mechanisms. There are potentially numerous interactions at each level, as well as bidirectional influences between levels. All these parameters may differ between different developmental stages or in different tissues of the body. (With kind permission from Springer Science and Business Media. *Metabolomics*, *Metabolomics – the way forward*, 1, 2005, p. 2, Goodacre, R., fig.a.)

the DNA), the proteome (the proteins produced under the control of mRNA), the metabolome (the products of proteins and other chemicals created by metabolism in the cell), and the phenome (the functioning of the cell within its environment in the body).

As shown in Figure 1.2, there are bidirectional arrows connecting these different levels (not a one-way arrow flowing from DNA to Function). One of the ways in which experiences may affect the expression of the genome is through the operation of control genes. Such control genes exist to control the operation of other genes by switching these other genes on or off (i.e., making genes either produce or stop producing mRNA). It now appears that such control genes may cause other genes to be switched off in response to changes in the internal and external environment. One remarkable example of such effects is shown by the observation that tweaking a rat's whiskers may cause changes in gene expression in the animal's sensory cortex (Mack & Mack, 1992). Similarly, when a songbird hears their species' song this experience may operate to change the expression of genes in the brain (Mello, Vicario, & Clayton, 1992). Thus, we need to accept that environmental effects may result in changes in the way genes are expressed. Such changes in gene expression may in turn result in long-lasting changes in the neural structures whose development is partly under genetic control (see Plomin, DeFries, McClearn, & Rutter, 1997, for more details).

8 *Understanding Developmental Cognitive Disorders*

In line with these findings from animals it has been shown that in human monozygotic (identical) twin pairs there are measurable differences in patterns of gene expression (differences in the genes that are active or being expressed). Furthermore, these differences in gene expression increase with age and tend to be greater for twin pairs who have lived apart for longer and who have experienced greater differences in lifestyle and health (Fraga et al., 2005). These effects clearly suggest that differences in experience produce different patterns of gene expression in people and that such differences may be responsible for differences in health and brain development that may have effects on behavior.

Summary

We hope that our discussion makes clear that the environment affects how our genetic makeup is expressed. The patterns of gene expression in cells will differ in different tissues and at different stages of development. The tissues most relevant for explaining differences in behavior are those in the nervous and endocrine (hormonal) systems. The most important point for the present argument is to appreciate that experiences may affect the processes involved in gene expression. Viewed in this way, the genome is not fixed in the way it operates throughout development. Rather, the genome receives signals from the environment that can turn genes on or off in different tissues of the body (including the brain). This means that differences in our experiences may well affect how genes that play a role in controlling brain development are expressed.

For most of this book we will be concentrating on explanations for developmental disorders that seek to relate observed impairments at the behavioral level to deficits at the cognitive level. We believe that such cognitive explanations are important and valid in their own right. A cognitive explanation of a disorder is essentially a functional explanation, couched in terms of how a particular skill is learned and performed, and in what ways this typical functioning is disturbed. Such an explanation is satisfying in its own right, and also has practical importance, in that it relates closely (though always indirectly) to how we can best assess and treat a disorder. This is not to say that biological levels of explanation are not also important. We will, where appropriate, cite evidence about the biological mechanisms underlying the cognitive level of explanation, particularly where such biological evidence places constraints on the types of cognitive explanation that are most viable. As has already been made clear from the brief account of research on conduct disorder above, there are two levels of biological mechanism that may be particularly relevant to the study of developmental cognitive disorders: genetic and brain mechanisms. We will consider very briefly the way in which these mechanisms are studied.

Genetic Mechanisms

There are two levels at which the genetic basis of a disorder can be studied. Population genetic studies examine the patterns of inheritance of a disorder across individuals.

Molecular genetic studies go beyond this and identify certain genes (DNA sequences) or gene markers that are associated with the development of a disorder. Both of these levels of analysis have been applied in the case of conduct disorder.

Population genetic studies relate variations in genetic association to degrees of similarity in the phenotype (observed characteristics). Basically, if a characteristic is inherited, people who are genetically similar to each other should also be similar to each other in that characteristic. One of the ways to get such evidence is from studies of twins. These studies make use of the fact that there are two different types of twin. Identical or monozygotic (MZ) twins develop from a single fertilized egg. Nonidentical (sometimes referred to as fraternal) or dizygotic (DZ) twins occur when two different fertilized eggs implant in the womb at the same time. MZ twins effectively share all their genetic material, whereas DZ twins will only share on average the same degree of genetic similarity to each other as any other pair of siblings. (DZ twins should, on average, share 50% of their segregating or polymorphic genes. These segregating genes are the coding sequences of DNA that differ between people and contribute to individual differences. Such segregating genes only account for a tiny proportion of our DNA: indeed it has been suggested that human beings share 98% of their genetic code with chimpanzees.) Twin studies often involve making comparisons between how frequently a disorder occurs in pairs of MZ and DZ twins. If both twins in a pair share the same condition, they are said to be concordant. Concordance rates should be higher in MZ, than DZ, twin pairs if genetic factors are important.

Concordance rates are only really useful when studying characteristics that are either present or absent. For example, if breast cancer were influenced by genetic factors, we would expect that the risk of pairs of MZ twins both contracting the disease would be higher than for pairs of DZ twins. However, as we shall see later in the book, for many cognitive disorders it is difficult to set precise cut-offs for whether a person has, or has not, got a disorder. This is because the disorders are best described as dimensional (so that individuals can have a disorder to varying degrees). Because of this we need a method of studying the degree of similarity between pairs of twins when the measures are quantitative dimensions rather than categories. Such a method was developed by DeFries and Fulker (1985). This method basically uses a form of regression equation to assess the influence of genetic factors on a characteristic. If genes are important in determining a continuous characteristic (such as height), MZ twins should be more similar to each other on that characteristic than DZ twins.

The degree of genetic influence on the development of a characteristic is expressed in terms of a heritability estimate. Heritability is concerned with quantifying the extent to which differences among people in a population reflect genetic differences. A heritability estimate of 0 would mean that genetic differences played no role in explaining the differences among people in a characteristic, while a heritability estimate of 1.0 would mean that genetic differences accounted entirely for the differences observed. In practice heritability estimates are usually intermediate in size but it is common for developmental disorders to show substantial heritability, meaning that genetic influences are important for their development. To return to the case of

10 Understanding Developmental Cognitive Disorders

conduct disorder, there is good evidence that genetic factors are important for its development. For example, Blonigen et al. (2005) reported a heritability estimate of approximately .50 for a measure of impulsive antisociality in a large twin sample, meaning that some 50% of the differences between people on this measure reflected genetic differences between people in the sample studied.

Molecular genetic studies try to identify the specific genes that may be responsible for the development of a disorder. Modern techniques allow the sequence of base pairs in an individual's DNA to be "read off" quite rapidly. The problem then becomes one of sifting the huge amount of data generated. It would not be appropriate to go into the details of these methods here. However, the basic approach is to try to identify DNA sequences that are shared by close relatives who both display a disorder but are not shared by other close relatives who do not have the disorder. Such studies involve sifting huge amounts of data and, rather than identifying specific genes, quite large DNA sequences (consisting of potentially many genes) may be identified. A group of genes that can be shown to correlate with the development of a complex quantitative trait (such as reading ability) is referred to as a quantitative trait locus (QTL). However, in some cases specific candidate genes have been identified that appear to be causally related to the development of a disorder. In the example of conduct disorder described above, one of the variants (alleles) of a gene coding for low activity of the monoamine oxidase inhibitor A (MAOA-L) appears to predispose an individual to display violent behavior (but only if they experience maltreatment in childhood).

The Causes of Development – Nature Working with Nurture

One of the oldest and most central debates in developmental psychology is about the role of genes (nature) and environment (nurture) as determinants of development. As we will see later in the book, there is overwhelming evidence that genetic factors are powerful influences on the origins of many developmental disorders. We take this conclusion to be established beyond any reasonable doubt. This is not the same as saying the disorders are innate, however.

Innate is defined in the *Shorter Oxford English Dictionary* as "Existing in a person (or organism) from birth ... inborn ... of qualities ... (especially mental) opposite of acquired ..." It is important to appreciate that the idea embodied in this definition is totally at variance with current thinking in genetics and developmental biology. The critical point is that genes contain information that serves to direct development, but all development takes place in an environment and information from the environment interacts with the genetic "blueprint" in complex ways. Development results from the interaction of genetic and environmental inputs – an idea referred to as epigenesis. Furthermore, according to the idea of "probabilistic epigenesis" (Gottlieb, 1992; Johnson, 1997), there may be bidirectional influences between different levels so that, for example, genes that help to specify aspects of physical development (including brain development) can in turn be reciprocally influenced by the structures they have helped to produce (see Figure 1.3). Similarly, and perhaps more obviously,

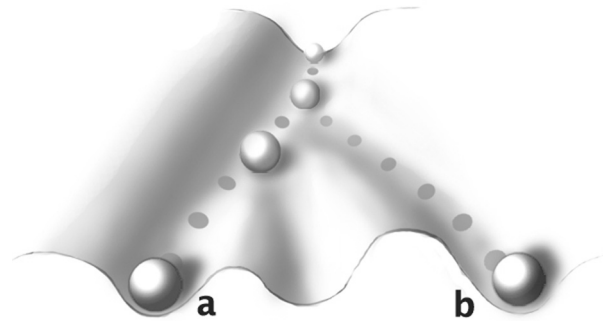


Figure 1.3 Waddington's epigenetic landscape is a metaphor for how gene regulation processes modulate development. Each of the marbles rolling down the hill represents a cell and different grooves in the landscape represent different trajectories that will result in different developmental courses and so different "end states" for a cell. Differences in the environment will play a role in determining the trajectory taken by a given cell. At a higher level we could think of the marbles representing whole organisms and again the end points of development will depend on both genetic and environmental influences.

learning (an influence from the environment) operates to modify structures in the brain that developed under genetic control and in turn may influence subsequent learning.

Development has to be seen as an extremely complex process that is characterized by change and interaction. All of the cognitive disorders we will consider in this book depend upon functional brain systems (brain systems that are defined by what they do) and it is simply not sensible to view these systems as arising directly and invariantly from information coded in the genes. In practice, performing any cognitive activity will depend upon one or more brain circuits, which comprise complex assemblies of many thousands of nerve cells communicating information between each other. Such brain circuits will develop under some degree of genetic influence but also as a product of learning from interactions with the environment.

Genes code for the production of proteins, which in turn have complex and at least partially indirect effects on the way physical structures such as the brain develop. Furthermore, as we have already noted, experiences may serve to switch on, or switch off, genes that are involved in controlling structural and functional aspects of brain development. In short, functional brain systems (brain circuits) develop as a result of complex interactions between genetic information and a range of environmental influences (where the environment includes many physical influences on development, such as temperature, nutrition, toxins, and radiation, as well as psychological experiences).

An acceptance that some aspect of development is under genetic influences does nothing to negate the importance of the environment. In relation to developmental disorders this can be illustrated by a well-known example. Phenylketonuria (PKU) is a genetic disorder that is controlled by a single gene. Children who inherit two such recessive alleles of this gene are unable to metabolize phenylalanine (an amino acid present in many foods) and this results in a build-up of this substance in the body

12 *Understanding Developmental Cognitive Disorders*

that damages the developing brain and causes general learning difficulties (mental retardation). However, PKU can be detected by a simple blood test (blood is taken in the heel prick test given to newborn babies) and provision of a special diet that is low in phenylalanine can prevent brain damage and the resulting learning difficulties from developing. A very clear discussion of the complex interplay between genetic and environmental influences on behavior is given by Rutter (2005b).

Brain Mechanisms

Genetic differences between people, in concert with environmental influences, determine the course of development, including development of the brain (epigenesis). In relation to developmental cognitive disorders it is likely that the problems we observe in different children will reflect both structural and functional differences in brain organization. In the last 20 years or so there has been an explosion of research concerned with understanding the relationships between brain, behavior, and cognition. Most of this research has focused on brain function, though it is also the case that some important work continues to examine the possible relationships between structural brain abnormalities and various forms of learning difficulties (Leonard, Eckert, Given, Berninger, & Eden, 2006)

Our ability to study the functional organization and operation of the brain while we are thinking has been transformed by the advent of brain imaging techniques. Positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) are two techniques that have been used to study the patterns of neural activation occurring during ongoing cognitive tasks. Both PET and fMRI detect changes in blood flow in specific regions of the brain that arise during the performance of a task. When a brain region does work, it requires metabolic energy, which in turn requires extra oxygen and thus extra blood flow. Both of these techniques provide evidence for fairly slow-acting changes in brain activity and usually depend on averaging measurements from a number of trials in an experiment. However, the techniques give quite precise information about localization in the brain. The other methodological wrinkle is that we need to have a “baseline” against which to measure any putative increase in activation in a specific task. This therefore involves subtracting the levels of activation seen in a specific task from levels of activation seen in a similar task, preferably in a task that involves everything apart from the one component of an experimental task that we are particularly interested in. So, for example, activation might be compared in a condition where a subject sees and silently reads a sequence of words, and in another condition where exactly the same words are presented as pictures to be silently named. Areas of the brain that show increases in activation in the reading condition, compared to the picture condition, presumably are somehow specifically involved in processing written words (orthographic processing) and translating from orthography (print) to phonology (speech sounds). Details of the subtraction methodology become complicated, but the point is that imaging studies always involve some sort of inference to be made based on a comparison between closely matched tasks.

Electroencephalography (EEG) and magnetoencephalography (MEG) are two techniques that give better temporal (time-based) information about patterns of brain activity but poorer information about the localization of activity. EEG involves attaching electrodes to the scalp and measuring differences in voltage between the electrodes and how these voltage differences change across time. The timing of these voltage changes, which reflect patterns of firing from large sets of neurons in the brain, can be measured with millisecond (0.001 s) accuracy. One particularly useful EEG technique is event-related potentials (ERPs). To measure ERPs, EEG recordings are taken in response to a particular stimulus (or set of stimuli) and the results are averaged over many trials to identify consistent patterns of activity. MEG is a methodologically superior technique to EEG that also measures changes in neural activity in the brain. MEG measures the magnetic fields produced by the electrical activity in the brain by using superconducting quantum interference devices (SQUIDS), which are housed in a helmet-like enclosure that fits around the head (see Plate 1). Like EEG, MEG yields quite precise information about the timing of neural responses to stimuli, but it gives relatively crude information about the localization of activity in the brain. It seems likely that MEG will become a very valuable technique for studying brain activity, and combining MEG with fMRI recordings in the same individual provides the possibility of getting both localization and temporal information about patterns of brain activity.

Separable Systems in the Mind – Modularity and Development

Subsequent chapters in this book will consider what we know about the nature, origins, and treatments for a variety of developmental cognitive disorders. The fact that there is a wide range of somewhat specific developmental disorders (some children have difficulties with language, while other children have difficulties with the control of movement, for example) supports the idea that the mind has different systems (or modules) that are responsible for different functions (language and motor control in the case just cited).

The idea that the mind is a modular system (a system composed of separable sub-systems) has a very long history that can be traced back at least as far as the ancient Greek philosophers (Arbib, Caplan, & Marshall, 1982). A slightly more recent, but now discredited, modular approach was represented in Gall's pseudo-scientific phrenology (see Figure 1.4). According to Gall the relative size of different brain regions (measured by feeling the shape of the skull!) could be used to infer characteristics of people, such as their "acquisitiveness" or "secretiveness." The idea of modularity has been brought to prominence in modern psychology by the work of Fodor (1983) and Marr (1983).

Studies of cognition in adults, and particularly studies in adult cognitive neuropsychology, have been dominated by an approach that sees the mind as a modular system. Cognitive neuropsychology seeks to develop theories about how the mind typically operates, by studying the disorders in mental (cognitive) processes that

14 Understanding Developmental Cognitive Disorders




Figure 1.4 A phrenological head showing areas labeled with their supposed functions.

arise as a result of brain damage (see Shallice, 1988). A modular view sees the mind as composed of separate systems or modules, just as we might think of our bodies as being composed of different systems such as the circulatory, respiratory, and digestive systems. An analogy to convey the concept of modularity can be given by considering a computer (see Figure 1.5). A desktop computer usually consists of a number of interconnected components, some of which are physically separate (the monitor, keyboard) while others may be housed in the same box (the processor, hard disk, CD drive, sound card, video card, etc.). Problems in such a system can be easily identified, and rectified, by isolating or swapping components. To take a trivial example, if the monitor does not work, this may be due to a number of components (the monitor itself, the cable connecting it to the computer, or perhaps the video card inside the computer that generates the signals to control the monitor). By testing each of these components sequentially we can gradually identify the component that is causing the fault in such a system (though often such a process can be time consuming and frustrating!).




Figure 1.5 A computer as an example of a modular system.

In some ways, the studies we describe later can be thought of as analogous to this process of finding a fault in a computer system. If, for example, children with dyslexia perform tasks requiring them to isolate individual sounds in words very poorly, but perform as well as other children on analogous tasks requiring them to isolate shapes in complex visual displays, we might infer that the brain systems dealing with speech sounds are impaired in children with dyslexia, while other brain systems dealing with the perception of complex visual patterns are intact. We will, however, spend a great deal of time showing how understanding disorders of the developing mind is a much more complicated process than locating a fault in a computer system.

Cognitive neuropsychology in adults has made enormous progress by adopting an approach that seeks to understand the effects of brain damage as arising from impairments to separable cognitive systems that can be damaged independently as a result of brain injury. At the simplest level, modularity simply amounts to the claim that the mind consists of separate subsystems. To take an obvious example, there are separate systems responsible for vision and hearing in the brain. Damage to the primary visual cortex (at the back of the head, in the occipital lobe) results in areas of blindness, while damage to the primary auditory cortex (at the side of the head, in the temporal lobe) results in difficulties in discriminating the frequency of sounds (Tramo, Shah, & Braid, 2002). In these cases no one would wish to argue with the proposition that separate brain systems are responsible for the senses of hearing and vision, and that it is possible to get impairments in vision, without impairments in hearing, and vice versa. This, in the parlance of cognitive neuropsychology, would be an example of a double dissociation: patients with damage to the primary visual cortex have problems with vision, but hear normally; patients with damage to the primary auditory cortex have problems with hearing, but see normally. Double

16 *Understanding Developmental Cognitive Disorders*

dissociations have often been interpreted as providing critical support for modularity: the existence of separable, neurally independent, systems.

This example has been chosen deliberately to be clear and noncontroversial. We will make the reasonable assumption for the time being that the two patients described showed massive deficits on the visual and auditory tasks, but that each was completely normal on the nonimpaired task (one patient had a severe visual impairment, but completely normal hearing; the other had a severe auditory impairment, but completely normal vision). In such cases evidence of this sort can be related to a variety of other evidence (e.g., that the primary visual cortex receives input from cells in the retina of the eye, and stimulation of the eye by a flash of light results in neural activity in the primary visual cortex) to support a theory that the visual system is functionally and neurally separable from the auditory system. However, such very clear cases are the exception, even in studies of adults following focal brain lesions, and such distinctions become much harder to make once we move on to consider “higher” cognitive processes such as memory. Furthermore, as we shall see later, in studies of children with cognitive disorders such clear patterns of selective impairment are quite unusual (and this is an interesting point in its own right, to which we will return).

In reality the logic and practice of seeking to establish the existence of separate cognitive systems by looking for double dissociations is both controversial and complex and has been debated extensively (e.g., Coltheart & Davies, 2003; Dunn & Kirsner, 1988; 2003; Gurd & Marshall, 2003; Jones, 1983; Van Orden, Pennington, & Stone, 2001). There are both logical and statistical issues at stake in the debate about this issue. Logically, it seems reasonable to conclude that any given pattern of double dissociation might in principle be open to a variety of theoretical interpretations. Claims about separable processes will always depend upon having a clear theory about the processes concerned and finding converging evidence to support the idea of their separability (as in the case of converging evidence for the role of the visual cortex in vision described above).

At another level there are also purely statistical or methodological issues about how we need to measure behavior in order to establish dissociations between tasks (which is a prerequisite for trying to infer that the tasks depend upon dissociable mechanisms). In a typical case, the process of establishing an impairment in one domain, but not another, amounts to identifying what Chapman and Chapman (1973) referred to as a differential deficit. As these authors pointed out, identifying differential deficits depends critically upon the statistical properties of the measures used. In particular, the greater the true score (or reliable) variance in a test, the easier it will be to show that a clinical group is impaired on that test. True score variance increases as the reliability and the variance (the range of scores) of a test increase. The reliability of a test refers to the extent to which measurement is subject to error. The variance in scores from a test will vary with the relative difficulty of the test for the sample of people it is used with; the variance in test scores will increase when tests are neither too hard (tendency toward a floor effect) nor too easy (tendency toward a ceiling effect). The statistical methods needed to identify differential deficits are well understood, though in practice these methods can be onerous and are rarely followed rigorously.




Figure 1.6 The hippocampus: a bilateral mid-brain structure that plays a critical role in memory function. Patient HM suffered bilateral damage to the hippocampus following brain surgery to relieve intractable epilepsy.

A much more severe limitation on the use of double dissociations comes from the sorts of inferences that we wish to draw from observing dissociations. In practice all we observe directly are dissociations on tasks (we observe that Patient A is bad on Task 1, but fine on Task 2, and that Patient B is bad on Task 2, but fine on Task 1). However, the interpretations that we are interested in involve theoretical systems (modules). Sometimes inferences from tasks to hypothetical processes seem straightforward enough. In the example given above (observing one patient fail to report flashes of light presented to certain parts of their visual field, and another patient with problems in identifying sounds) it seems safe enough to infer that they suffer impairments to the visual and auditory systems, respectively (and this assertion relates to a wide range of other evidence and theory). However, once we leave the study of “peripheral” sensory systems and move on to study “higher” mental processes, the inferences required become considerably more complex.

This difficulty of identifying separable higher mental processes can be illustrated with a justly famous example. Some adult patients have been identified who have particular difficulties with the immediate, verbatim recall of spoken lists of words while their memory for information from the past is preserved (KF, a patient studied by Warrington & Shallice, 1969). These problems have been interpreted as evidence for an impairment to a short-term memory system (a system holding information in conscious awareness for a second or two) with no impairment to a separate long-term memory store (a system that holds a vast amount of information for extended periods of time). In contrast, other patients (e.g., HM; Scoville & Milner, 1957; HM had bilateral damage to the hippocampus; see Figure 1.6) have been described who have great difficulties in remembering events from the past but are relatively good at the verbatim recall of lists of words, and this has been interpreted as evidence for an

impairment in a long-term memory system with an intact short-term memory store. However, note that in these cases the inferences from the observed tasks to hypothetical systems are much more indirect than in the earlier cases of sensory impairments.

Observations of these patients' deficits cannot establish that there are separable short- and long-term memory systems. There might be separate memory systems for immediate, as contrasted with delayed, memory but equally there may be other ways of explaining these patterns of impairments. For example, it might be that patients such as KF suffer problems in maintaining information in a phonological (sound-based) code (as may be required for immediate verbatim recall of arbitrary lists of words) while patients such as HM have problems with recalling semantic information (the meanings of events). This would still be evidence that memory is not a unitary process, but in this case there would be a distinction between the codes representing different classes of information, rather than a distinction in terms of the time period over which different memory stores hold information. The best interpretation of these dissociations need not be a matter of concern to us now. The point of importance is that the observation of a dissociation in performance on different tasks may be amenable to a variety of different theoretical interpretations (in terms of the underlying psychological systems, or modules, involved).

We now need to consider what precisely (or at least more precisely) we mean by modularity in psychology. Our example of a computer system is a useful intuitive starting point. A modern computer is a modular system in the sense that different components are specialized to perform different functions (such as the monitor, keyboard, CD drive, processor, sound card, and video card). Unfortunately, even with modern technologies such as brain imaging techniques it is not possible to identify separate components in the brain directly. Fodor (1983) set out to develop a theory about what cognitive modules are like. Fodor listed a number of features that he believed typically, but not necessarily, characterized modules. He suggested that modules tended to be "domain specific," meaning a given module only takes a restricted range of inputs (perhaps one module exists that "computes" the identity of letters, while another module "computes" the identity of spoken words). The first hypothetical module in this case would only take a restricted set of visual information as input (color and brightness would be irrelevant) while the second module would only take a restricted range of auditory information as input. This is analogous to the case of the sound card and video card in our computer example: the sound card plays no role in dealing with information going to the monitor, and the video card plays no role in dealing with audio signals going to the speakers.

Modules also, according to Fodor, display "informational encapsulation," which means that higher "conscious" levels of the system have limited access to the processes operating in the modules. Modules also tend to be "computationally autonomous," meaning that they do not share general-purpose resources such as attention. These three properties bring with them a benefit: modules, according to Fodor, operate quickly and effortlessly.

Fodor also suggested that modules were likely to be innate. (The suggestion that anything is "innate" is an idea that we have already rejected above.) We will not dwell on the details of Fodor's characterization of modules, which have been the

subject of extensive debate and criticism. The general point, however, is clear. Fodor was suggesting that there is a set of relatively independent, innate, fast-acting systems in the mind that carry on doing their work without the need for conscious attention or effort. Fodor also suggested that the evidence for modular organization was clearest in the case of input and output systems. He suggested that there were also “central systems,” such as those responsible for reasoning and problem solving, that probably were not modular. These central systems were, he suggested, very hard to understand, precisely because they were nonmodular. In fact, Fodor (2000, 2005) has gone on to argue that the view that much of mental life can be explained in terms of a large set of modules (“massive modularity”) is misguidedly optimistic. In his view the need for nonmodular systems underlying higher-level cognitive processes, and how such nonmodular systems can be conceptualized, remains a major challenge for theories in cognitive psychology.

In practice psychologists have generally been happy (perhaps too happy in Fodor’s view) to assume that the mind consists of many separate but highly interactive systems. Indeed without this assumption, or some version of it, the task of trying to understand the mind seems impossibly complex. As Fodor argued, this idea is generally easiest to grasp in relation to the most peripheral input or output systems (it seems obvious in the example above that different systems underlie vision and hearing) but becomes more difficult to pin down as we move to more “central” processes (how many different memory systems are there?). Psychologists typically seek converging evidence from different sources to support or refute theories about the number of separable modules that underlie particular skills or tasks. One particularly important source of evidence for separate processes (but not the only one) comes from demonstrations of separate impairments (people may be blind, but not deaf, and vice versa).

In summary, in studies of adults a common working assumption is that there are separate modules or subsystems in the mind that perform different functions. Theories are often expressed in terms of boxes and arrows diagrams, or more rarely as computer programmes that attempt to implement the processes represented in such separate boxes.

The Need to Relate Developmental Disorders to Patterns of Typical Development

How are we going to understand disorders of cognitive development? A number of different approaches have been adopted, one of which is to relate developmental disorders to patterns of impairment seen in adults following brain damage (Temple, 1997). In our view this approach is misguided because it is based on a view that the cognitive modules seen in adults are essentially innate. Our view, in contrast, would be that the only way to understand developmental disorders is to relate them to studies of typical development. If we want to understand the reading problems seen in children with developmental dyslexia we need to do this by relating the problems seen in these children to patterns of typical development. A theory of reading development needs to specify how typically developing children learn to read. A theory of developmental




Figure 1.7 Modular systems underlying the comprehension and production of spoken words in adults. Our ability to comprehend and produce spoken words appears to depend upon at least three separable systems dealing with the recognition (auditory input lexicon), production (phonological output lexicon), and meaning of words (semantic system).

dyslexia would then specify how and why the processes that are involved in typical reading development are impaired in children with developmental dyslexia.

If we accept arguments from adult neuropsychology and cognitive psychological theories of normal adult cognition that the mind is to some degree a modular system, we can then ask how such modules develop. One view, which Fodor favored, was that many modules are under strong genetic control, that is, we are born genetically programmed for the brain to develop certain specialized systems, given only certain rudimentary inputs. It seems plausible that this is true for some systems. Again, sensory systems such as vision seem plausible candidates for being under strong genetic control. However, there are plenty of modules postulated by psychologists that almost certainly are not under direct genetic control: the systems underlying reading and writing would be good examples. Reading and writing are very recently acquired skills in evolutionary terms and almost certainly reflect the fact that the brain can create modules by a process of learning. In fact, as Bishop (1997a) has pointed out, many of the properties of “modules” proposed by Fodor (such as speed and autonomy of operation) are plausibly seen as properties of cognitive systems that have benefited from extensive practice (think of skilled typing, or riding a bike, for example). No one is born with an innate “typing” module, but skilled typists can type at a remarkable speed, often with seemingly little cognitive effort.

Some have argued that a good starting point for studies of cognitive impairments in children is with a model of the adult “modular” system (Temple, 1997), one of the simplest things that can be said about development is that it involves change. Typically developing children show large changes in the course of acquiring certain skills such as language or reading and writing. To take a concrete example, there would be fairly wide agreement that in an adult we can distinguish at least three separate systems involved in dealing with individual spoken words: a system for recognizing the sound structure of spoken words, a system for saying words, and a semantic system that deals with the meanings of the words that we hear or need to say. This idea is illustrated in Figure 1.7.

The speech input lexicon holds representations of the words we know and can recognize when we hear them, the speech output lexicon holds representations of

words that we know in a form suitable for controlling how we say them, and the semantic system holds information about the meanings of words. A typically developing child learns to say their first words at around 12 months of age, and thereafter rapidly expands their expressive vocabulary over an extended period of development. Arguably, the typical 8-month-old has none of the ovals illustrated in Figure 1.7 in their mind, the typical 2-year-old has perhaps small-scale versions of all three systems, and the typical 10-year-old has something that is a very close approximation to the adult system. In our view then, it is not useful to start with an adult model and ask how it is impaired in a child who has severe problems in learning language. Rather, we would argue that developmental disorders require a developmental approach to understanding them, that is, we must start with theories and models of how, in a typically developing child, certain skills are acquired. The proper form of explanation for developmental disorders will be to specify how, and in what ways, the typical developmental path is disturbed. Although adult models may provide a very useful description of what typically arises as a result of development, such models say little about how systems develop. In this book we will take a developmental approach. For each of the disorders we consider, we will evaluate how current knowledge of the disorder may be interpreted in terms of theoretical accounts of typical development.

It will be clear from this that the challenges in understanding developmental disorders are considerable. In the case of adult disorders, we need to understand how damage to the brain results in a particular pattern of impairments. In terms of a cognitive model we need to relate the pattern of impairment to a static adult model of cognitive function. For developmental disorders the task is considerably more complex. This complexity arises from the fact that the developing cognitive systems we are studying change as children develop. A theoretical explanation of a cognitive disorder therefore needs to specify how impairments in a given process arise, and how this affects the development of other systems. Development involves change, and an abnormality of development means that the rate and pattern of change with age are modified. Furthermore, deficits in one area may have diverse knock-on effects on other aspects of development because different systems interact during development.

We might illustrate this complicated idea most simply by considering a sensory impairment such as deafness. Congenital deafness in childhood can have very damaging effects on oral language development, and speech skills will usually be severely impaired. Deaf children typically show grave deficits in a number of aspects of oral language development, including impairments in phonology (mastering the sound system of language) and syntax (grammar). These difficulties in a congenitally deaf child need to be understood in terms of how typical oral language development depends on our experience of hearing and producing speech. The effects of acquired deafness in adulthood are quite different, however. In this case, language skills that have been learnt previously remain intact although the comprehension of speech is obviously compromised by problems in hearing. This example conveys directly, we hope, how a problem early in development may be expected to have consequences that are both more diverse and possibly more severe than a corresponding problem that only occurs in adulthood after cognitive development is complete.

22 *Understanding Developmental Cognitive Disorders*

In summary, we believe some version of the idea of modularity (in a weak sense, meaning only that the mind consists of partially independent systems) is correct. However, we would reject explicitly the suggestion that most cognitive modules are “innate.” Rather, there is overwhelming evidence that cognitive systems depend upon extensive amounts of learning for their development. At a broader level we also believe that part of what occurs in development (and development here would include learning new skills, such as typing, in adulthood) is a gradual increase in “modularity” (Hulme & Snowling, 1992; Karmiloff-Smith, 1992). When a skill is in an early phase of development many systems (including, in Fodor’s terms, central systems such as thinking and problem solving) may be involved in its performance. As learning proceeds we suppose that neural circuits become established that enable the skill to be performed efficiently, effortlessly, and with less need for central control (think of learning to type or to ride a bike). Hence modules consist of established neural circuits that depend upon learning for their development. Such learning does not arise from a state of equipotentiality (from a brain where all structures are general purpose) and we accept that there are genetic and neural constraints on the development of the brain circuits (or modules) that underlie skilled cognitive tasks such as language, arithmetic, and motor skills.

The timing of development

Development involves change over time, and younger children are less accomplished in many different cognitive domains than older children. Developmental disorders are typically characterized by slow rates of development, either in a specific domain (specific learning difficulties such as dyslexia or mathematics disorder) or more generally across many domains (general learning difficulties or mental retardation).

A great deal of time has been spent discussing the extent to which different developmental disorders reflect delay (slow development) or deviance (abnormal development). Most of the disorders we will discuss seem best characterized in terms of a delay, at least early in development. For example, it is not that children with mathematics disorder are completely unable to perform arithmetic, it is just that their arithmetic performance tends to be slow and error prone, like much younger typically developing children. One of the main tasks in studying developmental disorders therefore is to explain the delays in development seen in such disorders: What process or processes are not working properly to result in such slow rates of development?

If development is slow it could be that problems will ensue because of critical periods in development. The idea of a critical period is that there may be a particular period of development when the child is prepared to learn a skill and that if learning does not occur during this period it will become difficult or impossible to compensate for this later in development. This idea is closely linked to notions of neural plasticity. Up to some point in development neural changes associated with particular forms of learning may occur easily but thereafter there is a gradual reduction in neural plasticity.

There is ample evidence from studies of animals for critical periods in development. One striking example is in studies of birdsong learning (for a review see

Brainerd & Doupe, 2002). Marler (1970), for example, showed that white-crowned male sparrows only learn their species-typical song adequately if they hear their own species' song during an early sensitive period between the age of 10 and 50 days. If learning is delayed beyond this point, mastery of the song will always remain impaired.

It is easy to see analogies between birds learning to sing their species-specific song and children learning to understand and produce their native language. It is probably no coincidence, therefore, that in developmental psychology the notion of a critical period has received most attention in relation to language learning. The idea of a critical period for learning language is closely allied to the idea that humans have an innate propensity to learn language. Lenneberg (1967) for example suggested that language learning needed to be completed by puberty, though others have suggested that the critical period for easy and complete mastery of language may be as early as before 5 years of age (Krashen, 1973). Pinker (1994) asserted that the acquisition of language is guaranteed for children up to the age of 6, is steadily compromised from then until shortly after puberty, and is rare thereafter. The evidence for a critical period (or, in a weaker form, a sensitive period) is not strong. Studies of children who have experienced severe deprivation in early childhood may in some cases show persistent problems with language, but in such cases it is hard to rule out the fact that these children had congenital abnormalities (Skuse, 1993). Other evidence comes from studies of second language learning, and it has been suggested that while adult second language learners may become perfectly proficient in their use of syntax, typically these people retain a persistent foreign accent in the second language, suggesting that the critical age for mastering the phonological system may be earlier than for syntax.

While the critical period hypothesis for language learning is probably too strong it does seem plausible that there may be a gradual decline with age in the ease with which we learn language and other cognitive skills. In studies of developmental cognitive disorders we are typically dealing with an impairment in the rate at which basic learning mechanisms operate. It may be in some cases that such a limitation in the rate of learning is further compromised by a gradual reduction in neural plasticity as children get older, though to date evidence for this idea is limited.

In summary, the most striking characteristic of the disorders we will consider in this book is that they typically involve slower rates of development in certain key areas. So, for example, a child with dyslexia may only learn to read slowly and with difficulty, while a child with mathematics disorder will show a similar pattern in relation to learning to perform arithmetic. However, the patterns of reading and arithmetic performance seen in such children usually resemble those seen in younger typically developing children: such patterns are described as delayed rather than deviant. It is possible that in some cases delays in the development of certain processes may result in deviant patterns of development in later developing processes. Such instances are probably rare and we will not discuss this issue further here. The most striking pattern that characterizes most developmental cognitive disorders is a delay in the rate at which particular skills develop.

Categorical versus Dimensional Views of Developmental Disorders

If most disorders we are dealing with reflect delays in development, this leads directly to questions of diagnosis. A delay in development represents a quantitative difference between children, and in this view it is natural to view disorders in dimensional terms. In this view children with mathematics disorder simply represent the bottom end of a distribution of children in terms of their mathematical skills. At first sight this view seems at odds with using categorical labels for these disorders. Instead of talking about mathematics disorder, perhaps we should refer to children with weak mathematical skills? The arguments about the usefulness of categorical labels in diagnosis are complex. Although we do subscribe to a dimensional view of the disorders we will be discussing in this book, we will typically use categorical labels to refer to affected children. This is partly just to aid communication: It is often easier to use a categorical label (children with dyslexia) than a dimensional expression (children with severe and specific problems in learning to read words).

The use of categorical labels for the extremes of continuously distributed differences among people is not confined to the area of developmental cognitive disorders. Think of medical conditions such as obesity or hypertension. There are large differences in weight and blood pressure between people in the general population (and incidentally these tend to correlate or be associated with each other). However, if weight or blood pressure becomes too high it may pose significant risk for other aspects of health. Nevertheless, exactly where we decide to put the cut-off between hypertension and “normal” blood pressure is to some extent arbitrary. The same is true of the disorders that we will be discussing in this book. We see these disorders as the extreme end of normal variations in skills in the population. However, diagnostic labels for disorders aid communication and can be useful in conveying the nature of the difficulties children experience and for guiding educational management and intervention policies.

Methods of Study in Developmental Cognitive Disorders

There are a number of methodological approaches that have been used to study developmental cognitive disorders. We will consider these approaches and their merits and weaknesses briefly.

Group versus case studies

In adult neuropsychology the detailed study of single cases has been particularly influential, arguably more influential than studies of groups of patients (Ellis & Young 1988; Shallice, 1988). In studies of developmental disorders both group and case studies have been used, but there is little doubt that group studies have been more important. It is therefore worth considering briefly the strengths and weaknesses of single case studies and group studies.

In adult cognitive neuropsychology the study of single patients has for many years been the dominant approach. The attraction of single case studies in adult neuropsychology is easy to understand and probably arises from the fact that some of the patterns of deficit seen in adults following brain damage can be extreme and remarkable. However, pure cases with theoretically interesting deficits do not walk into the clinic every day, so it becomes important to make the most of those rare cases that are available to study. There is no doubt that case studies of cognitive deficits following brain damage in adults have contributed powerfully to the development of theories of cognitive function (Shallice, 1988)

However, even if several patients exist with similar deficits, averaging the results from different patients may be dangerous because the group average may not be representative of any single patient in a group. This line of argument led Caramazza (1986) to argue that the study of single cases was the only valid basis for making theoretical claims in neuropsychology (cf. Shallice, 1988). This, to many people, seems to be going a little far; surely ten nearly identical patients are more convincing than just one? The basic idea, however, that carefully documented dissociations from single case studies provide strong evidence for the existence of separable systems has been widely accepted in adult neuropsychology.

In studies of developmental disorders single case studies have been used occasionally (e.g. Hulme & Snowling 1992; Pitchford, Funnell, de Haan, & Morgan, 2007; Temple & Marshall, 1983) but the dominant approach has certainly been to study groups of children. The reason for this, as Bishop (1997) has said so clearly, is that the aims of adult cognitive neuropsychology as compared to studies of developmental cognitive disorders have been quite different. Adult cognitive neuropsychology has been concerned predominantly with trying to make inferences about the structure of the mind from the patterns of impairment found after brain damage. This is captured very succinctly by the title of Tim Shallice's (1988) book *From Neuropsychology to Mental Structure*, which gives an excellent review of this area. The critical type of evidence for this enterprise comes from finding dissociations. Clear, theoretically interesting, dissociations only occur rarely in patients, but logically clear dissociations between different functions in well-documented cases are very persuasive. For these reasons single case studies have become the method of choice in adult cognitive neuropsychology.

Studies of cognitive impairments in children have not, however, been particularly concerned with identifying separate systems in the developing mind (though as we shall see in later chapters the occurrence of specific learning difficulties certainly provides evidence that we can identify separable systems that develop somewhat independently of each other). Rather, the major aim of those studying developmental disorders has been to understand the disorders themselves: their origin, developmental trajectories, and possible treatments. As Bishop (1997) has noted, a critical aim in studies of developmental disorders is to make generalizations about the patterns of deficit that characterize a particular disorder in order to identify its causes. In developmental dyslexia, for example (see Chapter 2), there is very strong evidence that the primary cause of most of these children's problems in learning to recognize printed words is a deficit in phonological (speech sound) skills. Such a conclusion

(a generalization about a group of children with a particular disorder) can only be reached by studying groups of children.

It is important to note that here we are making a claim based on an association or correlation: Reading problems in dyslexic children are associated with phonological difficulties. This association does not prove that the reading problems in dyslexia are caused by the phonological difficulties, though this is a plausible theory, as we will see later. In adult neuropsychology, researchers have been very wary of interpreting patterns of associated deficits as evidence for the organization of mental processes or structures. The fact that two deficits commonly occur in patients who have had strokes does not mean that these deficits are functionally related. The association may arise just because the brain regions that are involved in these two functions are close together and therefore liable to be damaged at the same time. One example of this comes from a disorder referred to as “Gerstmann syndrome,” which has been described in both adult (acquired) and developmental forms (see Shallice, 1988). Patients with acquired Gerstmann syndrome show a striking cluster of deficits including difficulties with arithmetic, spelling, right–left disorientation, and finger agnosia (problems in identifying the relative position of fingers by touch alone). Various theories were developed to account for the functional relationship between these different symptoms, but it is now generally accepted that the symptoms seen in Gerstmann syndrome cluster together merely because they all depend upon damage to anatomically adjacent brain systems in the left parietal cortex.

As we shall see in later chapters, studies of developmental disorders are replete with similar examples of associations leading theorists up blind allies. If we find that a certain disorder is associated with a particular cognitive deficit that does not mean we have found the cause of the disorder. Given that many associations occur, we will consider later in the chapter how we can try to determine which associated deficits may play a causal role in accounting for a disorder.

Longitudinal versus cross-sectional studies

Another critical methodological issue in studies of development is the distinction between cross-sectional and longitudinal studies. Cross-sectional studies look at children at one point in time. Most of the studies we will deal with in this book (whether case studies or group studies) are cross-sectional in design. Cross-sectional studies provide a “snap-shot” of development frozen at one point in development (or several, if different age groups are considered in a single study). Longitudinal studies in contrast assess the same children over a number of occasions and thus allow us to track how changes in one skill may relate to changes in another. Longitudinal studies have some very important advantages over cross-sectional studies but they are time consuming, expensive, and difficult to conduct. For these reasons longitudinal studies are usually only conducted when a number of cross-sectional studies have already identified some useful hypotheses that need to be tested in a longitudinal design.

One important advantage of longitudinal studies relates to how we interpret correlations (which is dealt with below). Most of the evidence we have about developmental disorders comes from correlations. So, for example, we may observe that a

group of children with dyslexia perform worse than typically developing children on several measures of speech (phonological) skills. This group difference amounts to a correlation; children with dyslexia also have poor phonological skills, and typically developing children have better reading skills and better phonological skills. Correlation does not demonstrate cause. Logically, there are three interpretations of such a correlation: (1) poor phonology causes poor reading, (2) poor reading causes poor phonology, or, most worryingly, (3) both things depend on something else we have not measured (e.g., general intelligence, motivation to complete the tests, language skills, the ability to attend).

Longitudinal studies allow us to assess correlations between measures taken at different points in time. Logically there is an asymmetry between correlations from measurements taken at different points in time. If a prior condition (phonology at Time 1) correlates with a later condition (reading at Time 2), the later condition cannot cause the earlier condition but clearly the earlier condition might cause the later condition. Longitudinal studies therefore help us to get a better handle on the direction of causation (but they are still open to objection (3) above; the only way round this is to conduct an intervention study).

The choice of comparison or control groups

In both cross-sectional and longitudinal studies (involving either groups or single cases) the aim is to identify the cognitive deficits that characterize a given disorder. The most widely used study design in this area is one where a group of children with a disorder is compared at one point in time to a group of children without the disorder. This type of design is called a case-control study in medicine. The question then becomes how to select the “control” children to compare to the clinical “cases.” Different comparison or control groups give us different types of information. The usual approach is to try to control for differences by matching the clinical cases to the control children on one or more variables. Given that on most tasks there are large increases in performance associated with increasing age, a common practice has been to select typically developing children of the same age as the clinical cases (a chronological age or CA control group). It would be common when doing this to try to match for other variables as well, such as the school(s) the children were attending, their gender, their scores on an IQ test, and other variables that might seem relevant to how well children would do on the experimental task being used. Comparison with a CA control group establishes if a clinical group has deficits on a task in relation to their age; if they do not we might conclude that the disorder is not associated with any difficulties on such a task. However, if a clinical group does differ from a CA control group on a given task (or set of tasks), this is really just the starting point for exploring whether such a deficit might be a plausible cause of a disorder. One obvious problem with such a finding is that the difference might be a product of the disorder rather than a cause (children with mathematics disorder might perform badly on a number judgment task simply as a consequence of their limited skills in, or experience of, mathematics for example).

One approach to reducing such problems of interpretation is to choose a younger typically developing control group of children who are matched for performance in the area of interest (perhaps reading ability as assessed by a well-standardized test of reading ability if we were studying children with dyslexia). This sort of comparison group is referred to as an ability-matched control group (in this case a reading ability or reading-age-matched group). This also is an informative comparison group and was first used to our knowledge in studies of children with spelling difficulties by Frank (1936).

One advantage of an ability-matched control group is that it eliminates absolute levels of performance on the task used to identify the clinical group (say, reading ability) as an explanation for differences on another task (say, speech perception). In other words, if children with dyslexia are worse on a measure of speech perception than younger control children whose absolute level of reading skill is the same as theirs, this difference cannot simply be a product of differences in reading skill. Conversely, a difference between children with dyslexia and a CA control group may always be explained away in terms of the difference being the product of a certain level of reading ability. While this is true, using an ability-matched control group usually means that there are large differences in chronological age (with the clinical group being older) and this may mean that deficits in the clinical group are unlikely to show up in such a comparison. If deficits do show up in an ability-matched design, this indicates that the clinical group are performing even more poorly than younger typically developing children, and this in turn suggests that the deficit is a severe one that we should consider seriously as a possible cause of the disorder. It is useful to have both CA-matched and ability-matched control groups to compare to a clinical group as each provides different information about the extent of difficulties shown by the clinical group.

Establishing the Causes of Developmental Disorders

The issue of how we can identify the causes of developmental disorders is really at the heart of this book, and we should consider this issue directly before we go further. A starting point for identifying a cause is to look for correlations. Though, as almost every introductory statistics text will tell you, correlation does not prove causation, the presence of a correlation is the usual starting point that makes us consider whether there is a causal relationship between two variables.

It may help to make this clear by looking at a well-established causal relationship from medicine. In the 1950s people asked whether smoking causes lung cancer. It was observed using case-control studies that there was an association (correlation) between smoking and lung cancer, with people who smoked being more likely to develop lung cancer than nonsmokers (Doll & Hill, 1950, 1954). However, a correlation is ambiguous because it might depend upon a third variable. It might be, for example, that a genetic difference between people causes both a greater propensity to smoke and a susceptibility to lung cancer. A very good way of expressing ideas about possible causal theories is in terms of path diagrams. Path diagrams, and the




Figure 1.8 Two alternative causal theories of the association between smoking and lung cancer. According to the first theory, smoking is a cause of lung cancer (and therefore persuading people not to smoke will reduce the risk of developing lung cancer). According to the second theory, smoking is not a cause of lung cancer (and therefore persuading people not to smoke will have no effect on the risk of developing lung cancer). The first theory is (probably) the correct one!

statistical techniques associated with them, arose early in the last century in studies of genetics by Sewall Wright (1920, 1921). In a path diagram single-headed arrows are used to express hypothetical causal relationships. These two alternative causal theories are expressed in path diagrams in Figure 1.8.

In fact, in the case of lung cancer further studies have filled in a detailed causal theory of how cigarette smoke causes lung cancer by inducing damage to the DNA in cells in the lining of the lung (Hecht, 1999). A simplified version of this causal theory is shown in the form of a path diagram in Figure 1.9. The details of this theory, and the evidence that supports it, need not detain us here. However, there is one aspect of this diagram that is useful in refining our discussion of what is a cause. The path diagram for lung cancer in Figure 1.9 shows a chain of causes, with certain constituents of tobacco smoke (including polycyclic aromatic hydrocarbons) leading to mutations in critical genes in cells of the lung. In terms of the theory shown here, cigarette smoke would be a *distal* cause of lung cancer while the gene mutations in the lung tissue produced by constituents of the smoke would be the *proximal* cause of lung cancer. It is clear that the terms proximal (immediate) and distal (distant) cause are relative terms that can change as theories develop. What appears to be the proximal cause of something in an early stage of theory development may become less proximal as further steps in a causal chain are uncovered and understood.

How, logically, do we move beyond correlations to prove what causes an illness or a disorder? This is a difficult question at the heart of scientific and philosophical inquiry that we will not try to solve here; see Pearl (2000) and Shipley (2000) for excellent, if technical, discussions of causality. We will, however, lay out some of the steps that we believe are useful in thinking about this critical issue.

Cause or causation is used to refer to a relationship between events. Informally, if every time a white billiard ball hits a red ball the red ball moves, we might say that being struck by the white ball causes the red ball to move. Traditionally, in philosophy it




Figure 1.9 A path diagram showing a more detailed causal theory of the processes that mediate the effects of smoking on the development of lung cancer.

was common to distinguish between necessary and sufficient causes. If something is a necessary and sufficient cause this implies certainty (every time A happens B will follow). However, causes could be necessary (we need A for B to happen) but not sufficient (sometimes when A happens B does not follow, perhaps because other factors are involved and are needed in addition to A to make B happen).

One problem with the use of the word cause is that in everyday speech “cause” often seems to be associated with this idea of a necessary and sufficient cause (i.e. with certainty of outcome). If smoking causes cancer anyone who has ever smoked should die of lung cancer. This is obviously not true. In fact current ideas about causality are framed in terms of probabilities, not certainties. Whenever we speak about causes in this book we will be talking probabilistically: Smoking we believe does cause lung cancer, but not in a deterministic fashion. Causes are things that increase the likelihood of an outcome: Smoking makes it more likely you will develop lung cancer. Furthermore, typically there are a number of causes operating to produce an outcome. Those different causes may sometimes operate independently (if smoking 10 cigarettes a day increases the chances of getting lung cancer by 10%, having a gene that makes you susceptible to lung cancer might also increase the chances of lung cancer by 10%; having both might then give a person a 20% increase in the likelihood of getting lung cancer). However, it is likely that often causes interact with each other, in which case, for example, having both smoking (10% increase in risk alone) and a susceptibility gene (10% increase in risk alone) might produce a 50% increase in risk of cancer. We saw earlier an example (in conduct disorder) where genetic and environmental causes interact in the genesis of the propensity to behave violently: Having a specific gene and being maltreated as a child both independently give a small increase in risk, but having both factors together gives a greatly increased risk of violent behavior.

Understanding that causes operate probabilistically means that a whole set of statistical techniques used to measure the association or correlation between variables is critical to assessing causes. Correlation and related techniques give us a way of expressing the strength of association between possible causes and their outcomes. In recent years huge strides have been made in philosophy and statistical methods for expressing and evaluating causal theories (see Pearl, 2000; Shipley, 2000). These advances give us a whole battery of conceptual and statistical tools to try to pin down causes. Sewall (Wright) who invented path analysis referred to it as “causal




Figure 1.10 A reproduction of one of Sewall Wright’s earliest path diagrams (Wright, 1920) showing the effects of genetic and environmental influences on the pattern of coat color seen in guinea pigs.

analysis.” Wright was a geneticist and his first paper to use path analysis (Wright, 1920) was concerned with the role of genetic and environmental influences on variations in coat color in guinea pigs. Wright’s first path diagram is reproduced in Figure 1.10.

There was no doubt for Wright (1920) that he was studying causal relationships: “In a broad sense the peculiarities of an individual are entirely determined by heredity and environment” (p. 328). Based on logic and theory he postulated genetic and environmental influences on the characteristics of animals born to parents who differed on the same characteristics. His insight was that it was useful to represent proposed causal relationships between things that could be measured in the form of a path diagram. In such a diagram measured variables (that stand as proxies for underlying hypothetical mechanisms) are connected by single-headed arrows to represent causal connections. Wright’s further insight was that, once such a diagram was constructed, measures of the correlations between variables could be used to estimate the strength of the proposed causal relationships, and that the strength of compound (indirect) causal paths could be calculated by simple rules (the strength of a compound path depends upon the product of the component paths).

Path diagrams are examples of what are now referred to in mathematics as directed graphs (directed graphs contain one-headed arrows, and arrows point from causes to consequences in path diagrams). Directed graphs, coupled with methods for translating them into quantifiable predictions using notions about probability distributions, give us a language for constructing and testing causal theories. Path diagrams are useful partly because they make such theories explicit and easy to understand and partly because they then link in to a powerful set of statistical techniques for testing our ideas about causes. (It should be noted that Morton and Frith’s causal

32 *Understanding Developmental Cognitive Disorders*

modeling framework for developmental disorders, with its use of arrows to link biological, cognitive, and behavioral levels of explanation, can be seen as direct application of Wright's (1920) idea of path analysis. The distinctive thing about their framework is that an explicit distinction is made between different levels of explanation.)

When is a correlation likely to represent a cause? Causes need to come before their consequences (after cannot cause before). This is called antecedence or the "logic of causal order" (Davis, 1985). If we notice a negative correlation between the amount of television watched as a college student and becoming President of the USA (people who become President did not watch much television as students), we might think that not watching television at college was a cause of later political success. But it could be the other way round (becoming President could not cause differences in television watching earlier in life!) (Davis, 1985).

The idea that causes must precede their consequences is why longitudinal studies are so important in trying to understand development and developmental disorders. The causes of a disorder should be observable before other symptoms of the disorder have developed, and correlations between possible causes of a disorder early in life and symptoms of the disorder measured later in life will help us to develop theories of what causes the disorder to develop. If phonological difficulties cause dyslexia, we need to show that such difficulties pre-date these children's reading difficulties (ideally we would want to show that phonological problems in young children, before they start to learn to read, predict the later development of reading problems when the children get to school).

We can try to home in on plausible causes by asking a large number of further questions once the issue of antecedence has been established (that the putative cause pre-dates the development of the disorder):

- 1 Universality? Does the deficit occur in all (or most) children with the disorder?
- 2 Power? Do variations in the severity of the deficit tend to correlate with the severity of the disorder?
- 3 Specificity? Is the deficit specific to the disorder, or is it also found in other diverse disorders (if a deficit occurs in many diverse disorders, it is implausible that it is a cause of any of them – because a cause should result in the same deficit in all affected children – unless such a cause is simply a contributor to many disorders that operates in conjunction with separate causes for each disorder)?
- 4 Theoretical plausibility? Finally, if a deficit does correlate strongly and selectively with one disorder we need to develop a theory of how a deficit in one skill (e.g., phonology) can lead to a deficit in the development of another skill (e.g., reading).

These arguments lay out how we can try to strengthen our belief in certain correlates of a disorder being causes of that disorder. Ultimately, however, no amount of correlational evidence can demonstrate a causal relationship. However, what can help a lot is having an explicit and detailed theory. Belief in causes depends a great deal on having a plausible mechanism that might explain an observed correlation (look back at the theory linking smoking to lung cancer). Saying we have identified

a “cause” of a disorder is always a theoretical statement and such statements are always provisional in science.

Can we ever prove that a certain deficit is the cause of a disorder? “Prove” is another very difficult word (just as difficult as cause). The traditional answer is that to prove that a variable is a cause we need to be able to manipulate it and observe a change in the outcome. This is the great insight (one of the great insights) of Sir Ronald Fisher (after whom the *F* ratio in statistics is named). Fisher (1926) described the principles of the randomized experiment in his book *The Design of Experiments*. Fisher worked in agriculture and we can use an agricultural example to illustrate the idea of a randomized experiment. Suppose we want to determine if fertilizer causes an increased crop yield from potato plants. We take a field and divide it into plots (say, 50 plots), we then at random apply fertilizer to half of those plots, wait for the crop to grow, and at harvest time weigh the potatoes from the 25 plots that received the fertilizer and the 25 plots that did not. There could of course be many differences between different parts of a field that might influence how well the potatoes would grow (differences in moisture and sunlight for example and perhaps differences between the potato seeds that were planted) but the process of randomization effectively serves to eliminate (or at least drastically reduce) the influence of such pre-existing differences on the outcome. The logic of random assignment is that it makes pre-existing differences between different plots or plants unlikely as explanations for the outcome of an experiment. The beauty of random assignment is that we do not need to know what these pre-existing differences are because by using a random process any such differences should all balance out.

To go back to the smoking example, and to dispense with all ethical considerations, if we wanted to prove that smoking causes cancer we could do the following experiment. Take a large group of people. Randomly assign them to two groups. Force one group to smoke cigarettes every day and ensure the other group did not smoke and lived in a smoke-free environment. We would wait several years and observe how many people in each group develop lung cancer. Random assignment of people to groups is essential here because it should eliminate pre-existing differences between the people that might influence the outcome. If we just asked for volunteers to smoke that would be hopeless, because only the reckless people (who for other reasons might be likely to contract cancer) would volunteer to be in the smoking group. Random assignment to groups is generally considered the best way for demonstrating causes in experimental sciences such as biology and psychology.

However, the notion of random assignment immediately runs into difficulties in relation to establishing the causes of cognitive disorders (as in many areas in biology and psychology). Given that the hypothetical causes of developmental cognitive disorders are typically properties of the child (their genes and brain systems that have developed under genetic influence), we cannot use random assignment to establish a causal link. By definition, we cannot randomly assign some children to have a certain genetic makeup, or to have a particular highly circumscribed cognitive deficit that we might believe is the cause of a developmental disorder. Such a procedure is logically impossible to implement (which may be just as well because, if it could be done, it would certainly be unethical). In this sense our room for maneuver in testing

the causes of developmental disorders is severely limited. There is, however, one version of a randomized experimental design for testing causes that we can use – it involves intervening to try to remediate a cognitive deficit.

Intervention studies as a way of establishing the causes of a disorder

Let us suppose that we have developed a causal theory about a cognitive deficit that we believe is a critical cause of the symptoms displayed in a disorder. (Much of the rest of this book is devoted to sifting evidence relevant to assessing the plausibility of theories of this type). If a potential cognitive deficit is identified we can then try to develop an intervention (usually a form of education or training) to cure (or at least ameliorate) the cognitive deficit. We can then randomly assign children diagnosed as having the disorder to either receive the intervention or to receive another form of intervention that targets another area of cognitive function. If such an intervention study “works” in the sense of producing improvements in the putative cause of the disorder coupled with a reduction in the frequency or severity of the symptoms of disorder, we have found good evidence for a cause (at a cognitive level) of the disorder.

This might seem a little circumspect, which is intended. Evaluating whether such an intervention works is a complicated process. There are two important aspects to evaluating if such an intervention has worked. First, the intervention should produce an improvement in the symptoms of the disorder. Second, ideally we should be able to relate the degree to which the intervention has improved the symptoms of the disorder to the degree to which the targeted cognitive deficit has improved. (If we believe that the reading difficulties of children with dyslexia depend upon a phonological deficit, then we should be able to improve these children’s reading skills by training phonology, and the extent to which their reading skills improve should relate to how much of an improvement in phonological skills the training program has brought about.) These requirements are an ideal, and in practice there are many complications in conducting and evaluating such intervention studies. Nevertheless such intervention studies are potentially of enormous theoretical importance; they get us as close to establishing the causes of a disorder as we can. Such studies are also of great practical importance because if they work they lead directly to practical recommendations about how to treat or prevent the development of a disorder.

A final word on causes: The importance of theories

Even in the case of a successful intervention study (or set of such studies) we always need to be careful about claiming to have “proved” a cause. Causes ultimately depend upon a theory: a model of how things operate. These theories depend upon processes that can never be directly observed but must instead be inferred from observations, and the way we make our observations and analyze the data from them typically involves numerous untested assumptions. According to Karl Popper (1980), in science we never prove something. What we do is develop theories to explain our observations. Any good theory is testable, and can therefore be disproved

or refuted by further observations. This is the spirit in which we should approach the testing of theories about the causes of developmental disorders. We observe the disorder and measure things about it. We formulate theories about the causes of the disorder in the ways we have outlined. Good theories will give an explicit statement about how and why a disorder develops. Further studies of the disorder will in turn test predictions from the theory and likely refine, alter, and sometimes even refute it. In Popper's terms all we ever have are conjectures and refutations (theories and problems for these theories). No amount of positive evidence (supported predictions) ever proves a theory. But negative evidence (unsupported predictions) may lead to a theory being abandoned or changed. In practice no single study or observation is ever water-tight, which is why in science we place such importance on the notion of replication. If a whole series of studies carried out in slightly different ways, by different people, fail to uphold a critical prediction of a theory, that theory is weakened and may be modified, or eventually be abandoned, in the light of accumulated negative evidence. Science is a sifting of evidence in which we try to get closer and closer to an approximation of the truth about how things work. The "truth" here is always an abstract model or theory of how things operate. Based on an explicit theory we accumulate and sift through evidence trying to evaluate in diverse ways the adequacy of the theory. This sifting of evidence may not be too dissimilar to the weighing up of evidence in a court of law to arrive at a judgment about the probable guilt or innocence of a defendant (Rapport and Wright, 1963).

Comorbidity and Separating Causes from Correlates

We need to be clear that many problems that occur commonly in a cognitive disorder may not be the cause of that disorder. It is now well established that different developmental disorders often tend to co-occur in the same child – this is referred to as "comorbidity" (Angold et al., Costello, & Erkanli, 1999). Caron and Rutter (1991) showed that comorbidity between many developmental disorders occurs much more frequently than expected by chance, given the rate of occurrence of the different disorders in the population. Getting accurate estimates of the true rates of comorbidity is difficult because it depends upon having representative samples of the population (such samples require epidemiological studies in which we assess many children who are selected to a truly representative sample of children in the population at large). So, for example, if we believe that motor impairments tend to co-occur with reading problems, we could test all the children referred for reading problems to a clinic to assess both their reading and motor skills. This would not allow us to get an accurate assessment of true comorbidity, however, because the children referred to the clinic are unlikely to be truly representative of all children with reading problems (perhaps having both a reading and motor disorder makes it more likely that the child will be referred to the clinic).

Comorbidities may arise for different reasons in different disorders. In some disorders, it is plausible to argue that one disorder may cause another. So, for example, having attention deficit hyperactivity disorder (ADHD) early in life may put children

at risk of developing conduct disorder later in life. In this example, perhaps a fundamental problem in inhibiting or regulating behavior shows different manifestations at different stages of development. In other cases, such as the comorbidity between reading difficulties and motor impairments, it seems unlikely that there is any direct causal link at the cognitive level between these two disorders; both problems may simply reflect the fact that brain development has gone awry with diverse effects on cognitive development. So, in this case, we would argue that problems with balance and motor coordination are unlikely to tap a cause of problems in learning to read. We will say much more about comorbidities between different disorders in the chapters that follow and return to the issue in Chapter 9.

Summary and Conclusions

In this book we will review what we know about a range of developmental disorders of cognition. We will consider a wide range of disorders including reading disorders, language disorders, arithmetic disorders, motor disorders, and autism spectrum disorders. This chapter has outlined a number of key conceptual issues that lie at the heart of studies of developmental cognitive disorders. We have argued that such disorders can only be understood in the context of a developmental theory of how the cognitive processes concerned typically operate, and by inference how those developmental processes are delayed or disordered in some children. As we shall see, most developmental disorders seem best characterized in terms of delays to typical developmental processes. This leads us to see such disorders as dimensional: The children identified are simply at the bottom end of a continuum of normal variations in the population. Nevertheless, diagnostic labels can be useful in communicating the form of difficulties experienced by different groups of children. Finally, the process of trying to understand the causes of developmental disorders is highly complex and depends critically upon having explicit theories of the nature of each disorder and the causal processes operating at different levels to generate the behavioral profile seen. Such causal theories can be tested in a variety of ways, but two of the most powerful forms of evidence come from longitudinal and intervention studies. We have introduced the idea of path diagrams as ways of representing causal theories, and we will use such diagrams in different chapters to represent different theories about the origins of the disorders we discuss.