1

Introduction

What Is the Psychology of Attention About?

One of William James's most famous statements was that "everyone knows what attention is" (James, 1890/1950), echoing an earlier claim by Munsell: "On attention itself, it is needless to discourse at length; its nature and conditions are familiar to every thoughtful student" (Munsell, 1873, p. 11). The present book takes a more empirical and more skeptical tack, assuming instead that no one knows what attention is, and that that there may even not be an "it" there to be known about (although of course there might be). The rationale for rejecting James's dictum will be discussed shortly. First, however, it should be acknowledged that there is certainly something to what James said. The word "attention" is commonly encountered in ordinary language, and people seem to understand each other well enough when they use it. This suggests that they must share some important notions and experiences in common. It is not unusual, of course, for an area of research to have a name that also figures in common parlance; for example, fields of study are devoted to bacteria and sedimentary rocks. However, although people talk about bacteria or about sedimentary rocks, they do not often have articulated theories or strong convictions that were not arrived at, directly or indirectly, from scientists and science teachers. In contrast, people talk about attention with great familiarity and confidence. They speak of it as something whose existence is a brute fact of their daily experience and therefore something about which they know a great deal, with no debt to attention researchers.

It is worth reflecting briefly on what most people think they know about attention as a preliminary to approaching the topic scientifically. Two primary themes or aspects characterize the phenomena people allude to with the term attention: selectivity and capacity limitation. Selectivity is apparent in a number of undeniable facts about human experience and behavior. One is that conscious perception is always selective. Everyone seems to agree that, at any given moment, their awareness encompasses only a tiny proportion of the stimuli impinging on their sensory systems. The second fact is that this selectivity holds not only of conscious states of mind, but also of the impact of stimuli on behavior. Whether we are walking, driving, playing tennis, or choosing what book to pick up, the stream of behavior we produce reflects only a small subset of the sensory stimulation, and typically only a subset of the stimuli that could potentially guide the same types of behaviors (e.g., the page contains other text that we could read; other paths are available for us to walk along). In short, the mind is continually assigning priority to some sensory information over others, and this selection process makes a profound difference for both conscious experience and behavior.

The second phenomenon to which casual usage of "attention" alludes is our limited ability to carry out various mental operations at the same time. Two activities that a person can easily carry out one at a time often pose tremendous problems when attempted simultaneously (e.g., listening to the radio and reading a book), even when these activities are in no way physically incompatible. This applies to many kinds of mental activities, including analyzing new perceptual input (as in listening to the radio and reading a book), thinking, remembering, and planning motor activity.

In the commonsense metaphysics implied by our ordinary talk about mental life (what some philosophers like to call folk psychology), selectivity and processing limitations are both described as attentional because they are really just the flip side of the same coin. The coin, of course, is attention itself. Folk psychology postulates a kind of substance or process (attention) that can be devoted (paid) to stimuli or withheld from them. Whether or not attention is allocated to a stimulus is usually thought to depend on a voluntary act of will; in the metaphysics of folk psychology, this ultimately depends on a choice made by the self. Sometimes, however, attention is directed or grabbed without any voluntary choice having taken place, even against strong wishes to the contrary; this is the phenomenon of distraction. Selectivity arises because awareness of a stimulus, and subsequent memory for it, is usually assumed to occur only when that stimulus is attended, voluntarily or involuntarily. The available quantity of attention is assumed to be finite, and this finiteness is supposed to account for capacity limitations; paying attention to one thing means having less attention available to pay to other things. Attention, according to folk psychology, can be devoted not only to stimuli, but also to activities, tasks, and behaviors (as in "pay attention to your driving"). Allocating more attention to any given task enhances performance, it is thought, although for certain highly automatic activities that do not require attention, such as breathing or tying one's shoes, too much attention may even be harmful.

A variety of other notions play a role in folk psychology talk about attention, and some of these issues go beyond selectivity and limited capacity. One is effort—the sometimes aversive experience of exertion that generally accompanies sustained mental activities; in folk psychology, this is often equated with the attention demanded by the task. Another example is perceptual set: the tendency for our perception of a stimulus to change depending on whatever expectations we may have about various aspects of the stimulus. It is often said that people see or hear what they expect to see or hear, and this selectivity is usually linked with attention.

Thus, our commonsense metaphysics of mental life points out, and in a very loose way might be said to try to explain, a number of phenomena: selectivity of perception, voluntary control over this selectivity, and capacity limits in mental functioning that cannot be attributed to mere limitations in our sensory or motor systems. These are the core phenomena addressed by attention research, and all will be discussed in this book. Unfortunately, the metaphysics of mental life that goes along with the commonsense use of the term attention is so entrenched and so compelling that whenever one uses the word one finds these notions intruding on one's thinking. This has several consequences, some innocuous and others pernicious. For example, many psychologists use "attention" as a theoretical construct in accounting for some range of phenomena, saying, for example, "I will measure the rate at which attention is shifted toward a location in the visual field" with regard to a study of visual cueing. This can be taken as a mere shorthand for "I will postulate a mechanism that can be allocated to certain input channels defined by location, and I assume for the sake of discussion that it probably has something to do with what ordinary people refer to as attention," leaving open exactly how close this linkage may be. Less innocuous are the many cases in which the notion of attention smuggles its way into the description of experimental manipulations (e.g., attention was reduced by having the subject remember seven digits) and observations (spare attentional capacity was measured by response times), obscuring the fact that these manipulations could have many different interpretations that ought to be considered, some of which would satisfy few or none of the assumptions contained in the ordinary conceptions of attention.

The dangers of taking substantive words from ordinary language and assuming a corresponding entity have been noticed for a long time, of course; philosophers at least as far back as Bacon (1620/1960) have warned against assuming that where there is a word there must be a thing (this is often called reification). Closer at hand, one of the pioneers of modern attention research, the late Donald Broadbent (1982), emphasized the dangers in the uncritical use of "attention." However, like all vices, reification is more easily deplored than prevented, and examples such as the ones mentioned in the preceding paragraph can be found throughout the present-day literature on attention (including some articles written by the author).

To try to avoid these pitfalls, the word is used sparingly in this book. When the goal is to uncover generalizations that are fairly close to the data, a special effort will be made to prevent assumptions imported from common sense from obscuring the description of phenomena or findings. Banishing the word totally would require cramped language, however, so an intermediate strategy will be followed. In chapters 2 through 4 that focus on empirical phenomena, the use of "attention" will be restricted to describing the field of study or the instructions given to a subject, as in "In a study of selective attention . . . " or "Subjects were told to attend to the tones played in the left ear . . . " In the same spirit, stimuli will sometimes be described as attended or unattended. "Attention" will not,

however, be used to refer to a putative internal process or mechanism (as in "Is attention necessary for memory?" or "The subject's attention was diverted with a secondary task."). In the fifth and sixth chapters the role of attention as a theoretical construct will be considered explicitly and critically, and a framework postulating several distinct attentional mechanisms will emerge in the remaining chapters.

Approaches to the Analysis of Attention

Between 1880 and 1920 a great deal was written about the psychology of attention by some very insightful writers such as Oswald Kulpe, Edward Titchener, and William James.[1] These psychologists inquired about the nature and significance of attention using terms and concepts borrowed directly and unapologetically from commonsense psychology—the self, the will, conscious perceptions and sensations, and attention itself. Their method of attack involved a mixture of introspection and, to a lesser degree, formal experimentation. The questions these writers asked make good sense in the context of this framework. Consider a few examples. How does attention to a stimulus change the conscious perception of the stimulus? Does it make the stimulus more intense? Does it make it more clear? Are clarity and intensity really different things? These four questions probably generated the most vigorous controversies in the early literature (see, e.g., Newhall, 1921; Titchener, 1908; Kulpe, 1902; Munsterberg, 1894). Divided attention was also a topic of interest, starting with the basic question of whether it is possible to attend to several inputs at once (see, e.g., James, 1890/1950, p. 405). Another topic of interest was how the speed of responses to a stimulus depends on whether a person chooses to attend to the stimulus or to the response (e.g., Cattell, 1893/1947).

The fact that many disagreements arose in this literature and that introspective methods were unable to resolve them are quite notorious. This should not obscure the fact that many of these authors succeeded in describing the texture of subjective mental life with great subtlety and clarity, and actually shared many important points of agreement. On the other hand, when it came to uncovering mechanisms and processes—asking about the workings of attention—the writers lacked concepts and terms with which to begin. For example, on the basic phenomenon of attentional selectivity, early writers strove continuously to reduce somehow the inner act of selection to overt acts of selection such as eye movements and postural adjustments (see Smith, 1969 for a review). However, most thinkers of the time, including James and Kulpe (along with Helmholtz), recognized that inner acts of selection simply did not have to be accompanied by any overt signs whatever. Having concluded that, they lacked any concepts with which to press on and ask more about what *was* going on. Consider, for example, James's description of the operations of selection: "[there is an] anticipatory preparation from within of the ideational centres concerned with the object to which the attention is paid." The language is evocative but the content is unsatisfying. Paschal (1941) suggests that early attention theorists developed a certain malaise on account of their difficulties in moving beyond these kinds of conceptualizations.

Before long, behaviorism and learning theory came to dominate American psychology for some decades, relegating the area of attention to some degree of obscurity. Behaviorists were seemingly uncomfortable not only with introspective methods, but also with the existence and demonstrable potency of selective attention itself. Selectivity is a problem for anyone who wishes to predict behavior purely on the basis of the objective stimulation making up the animal's learning history. For all these reasons, the behaviorist period was marked by a decline in psychological work in attention, although as Lovie (1983) points out, the disappearance was less complete than is sometimes claimed.

The Information-Processing Approach

This book discusses attention within the framework of the contemporary, information-processing approach. This framework had its origins in research carried out in England during and after the Second World War. It views the mind as an information-processing system whose function can, it is hoped, be revealed most clearly by making systematic observations of the fine-grained features of human behavior in laboratory situations. The ultimate goal is to work out some account of how the mind processes information, an account that is mechanistic but posed at the functional level. It is sometimes mistakenly supposed that this tradition inherently

depends on some likening of the mind to a particular kind of computer, especially the conventional digital computer. In fact, human information-processing research may provide the clearest evidence about how different the mind and the digital computer are in their underlying modes of function. The core idea of the information-processing approach is to analyze the mind in terms of different subsystems that form, retain, and transmit representations of the world. The nature of these subsystems, the kinds of transformations they carry out, and the temporal relationship and relative discreteness (or nondiscreteness) of their activities are all viewed as facts to be discovered, not assumed at the outset. Obviously, one cannot say in advance whether or not the mind can be successfully analyzed in these terms. Many optimistic forecasts have been offered; other writers began pronouncing the demise of this approach at least twenty years ago when it was barely under way (e.g., Newell, 1973; Anderson, 1987).

There is probably no good a priori reason to assume that a system as complex as the brain must allow of a functional description that is illuminating and reasonably accurate. After all, it seems conceivable that the only accurate mechanistic description might be one of neurons and their interactions, with mental activity being too nebulous to be characterized. Even if a true functional analysis is potentially available, we have no guarantee that we can arrive at this analysis by studying the input-output behavior of the system, or, for that matter, by studying its physical components. Trying to decide what can in principle be learned from any particular type of research is probably futile; the success or failure of information-processing psychology can be assessed only on the basis of the insights that do or do not emerge from the research it spawns.

This book will try to convince the reader that behavioral experiments can provide important information about attentional phenomena and many illuminating hints about the basic architecture of the mind. A fuller analysis of this architecture may well require recourse to neural as well as psychological levels of analysis, but in that case it seems unlikely that neural observations devoid of careful functional analysis will reveal much that is interesting about cognition. The reader will have to make his or her own judgments about how convincing and illuminating the functional analysis is proving to be.

What kind of data are to be used in analyzing human information-processing machinery at a functional level? As noted, early attention theorists relied on introspection. Introspection is now widely agreed to be an inadequate basis for arriving at insights into how mental processes work. It should not be dismissed entirely, however. Few writers today would claim that introspection can tell us much about processes operating at very fine time scales—say, a few hundreds milliseconds or less. For example, if you try to recall your telephone number, you probably find that the number pops into your head without your having any way to decompose this experience into finer elements. On the other hand, if you plan a strategy in chess, you probably find that you consider a number of alternative moves and countermoves one after another in a sequence. It seems dogmatic and unreasonable to claim that these latter introspections have no relationship to the events that culminate in the choice of a chess move (cf. Ericsson and Simon, 1980). Then in many intermediate cases people seem to want to claim at least some direct access to their mental activities. An example is the weighing of different factors in choosing which car to buy; people often have opinions about the role of different factors in a decision, but the validity of these reports is debatable (Nisbett and Wilson, 1977; Kraut and Lewis, 1982).

Where does attention lie along this spectrum? Many of the important questions about attention clearly fall into the first category: processes that occur rapidly and seem opaque to the individual. This does not imply that introspections about these phenomena are uninteresting. After all, if different individuals tend to agree about a certain aspect of experience—and the writings of classic attention theorists contain numerous points of agreement—then one could reasonably expect a complete understanding of mental architecture to explain *why* they should agree, even if this understanding does not give any credence to the actual assertions that people agree upon.

Consider a concrete example. Writers in the eighteenth and nineteenth centuries were interested in whether people could divide attention among more than one object, and their introspections on the topic differed tremendously. Some, such as Dugald Stewart (1792/1971), insisted that it was impossible. Others, like James, said it was not possible unless an entire collection of things was apprehended as a single object. Munsell (1873), by contrast, had no difficulty reaching a completely different answer: "Can a man attend to more than one thing at a time? To this the answer must be decidedly, yes!" (p. 12). Porter (1868) agreed, arguing that our ability to compare implies simultaneous consciousness of two things. The contemporary approach as represented in this book begins with the assumption that in the first place there may or may not be anything called attention to be divided, and that if such an entity does exist, a person's introspections would not necessarily provide accurate information about whether it could be divided. However, to the degree people are consistent in what they report about their experiences when trying to divide their attention, one would hope that a good psychological theory of attention would have something to say about this. For example, in a whole-report task, in which people try to report as many letters or digits as possible from a briefly flashed array (discussed in chapter 3), observers can usually report about four or five items. Almost invariably, though, they report that they "saw" many more (e.g., Sperling, 1960). A good psychology of attention does not have to include "seeing" as a theoretical term and endorse the observers' claims that they see all the letters, but it ought to have something to say about why people so often seem tempted to make that claim.

Theory Testing

Broadly speaking, how are theories about information-processing mechanisms to be tested using behavioral data? What kinds of inferences are to be trusted? One rather idealized approach to theory testing is sometimes termed strong inference (Platt, 1964). The essence of this strategy comes from Bacon (1620/1960), who said, "It is granted only to proceed at first by negatives, and at last to end in affirmatives after exclusion has been exhausted." The strategy consists of trying to test broad classes of models or very general claims about a phenomenon, aiming for critical tests; that is, tests that could potentially rule out the alternatives, ultimately leaving only the correct account standing. In a well-known article published in *Science* in 1964, Platt eloquently described the success of this approach in many areas of biology and physics. Broadbent's seminal book on attention (1958) advocated exactly the same strategy, offering an illuminating analogy between theory testing and the game Twenty

Questions. In this game one player thinks about something, and the other players try to figure out what he or she is thinking about. The other players are allowed to ask the first player only a series of yes-or-no questions. To progress as efficiently as possible in this game, one has to rule out large classes of possibilities at a time, and this is best done by asking very general questions. Posing highly specific questions at the outset (are you thinking about a Zebra?) is inefficient, since there are so many possibilities at that level of specificity. Similarly, both Broadbent and Platt argued that optimum scientific reasoning involves asking progressively more fine-grained questions, making certain to use tests capable of ruling out competing alternatives. Better known than Platt's elimination-oriented strategy is Popper's dictum that a theory is supported to the extent that it successfully predicts the outcome of experiments or observations that could potentially have falsified the theory. According to Popper (1959), it is the hallmark of a sound scientific theory that it should undergo the risk of falsification; the more risks, the more credible is the theory that survives.

As many philosophers and scientists have pointed out, the ideals of critical experiments and falsificationism are reached only rarely. There are various reasons for this. Often one cannot enumerate all the possible hypotheses that might be considered making the idealized strong inference strategy impossible. Furthermore, empirical tests invariably rely on ancillary assumptions of various kinds (Lakatos, 1978). Nonetheless, the Popperian and strong inference strategies offers a useful standard for evaluating experiments and inferences based on experimental results. One should probably expect, however, that in areas as treacherous psychology, critical experiments will be relatively rare. More often, after a series of nearly critical experiments, one theory may account for the data in a fairly straightforward way; competing accounts can survive too, but only by shouldering a large number of unreasonable and ad hoc modifications.[2] In this way, truth may be reached in the limit; from a Bayesian perspective, the posterior probability of the theory may continue to increase.

Another useful concept (perhaps especially in psychological research) is the idea of converging measures (Garner, Hake, and Eriksen, 1956), alternative tests that pose the same questions but rely on different kinds

of assumptions. The rationale is simple enough: when one seriously doubts the validity of the assumptions that underlie one test or another, one gains leverage by using tests that make assumptions that are independent of each other. Another intuitively natural observation is that experiments that make predictions for experiments whose outcomes have many degrees of freedom provide stronger support than experiments with only a few relevant outcomes. This notion can be readily justified in a Bayesian framework (Howson and Urbach, 1989; Jefferys and Berger, 1992). Unfortunately, this feature is not present in experimental psychology as often as one might hope. An experiment with three possible outcomes (e.g., $A > B$, $A = B$, $A < B$) supports a theory that correctly predicts the outcome, especially when competing theories make a different prediction. However, it cannot increase its credibility as much as an experiment with 20,000 possible outcomes, where the theory predicts the correct one, or, as sometimes occurs in physical sciences, predictions prove accurate to six decimal places or better.

These generalizations about theory testing may strike the reader as fairly banal and not worth belaboring. Nonetheless, a surprisingly large amount of research in experimental psychology, especially in leading journals of psychological theory, pays little heed to these principles. Instead, detailed and complex formal models are constructed and applied (usually with simulations) to fit the data. Typically, investigators have broad theoretical claims in mind; as a test of these claims, they construct a model that instantiates them, along with enough extra machinery to "predict" specific observations (e.g., a reaction time, a percentage correct). This machinery can be quite extensive. Papers of this kind typically present figures that compare "observations" and "predictions," usually with a stunningly close fit. Apparently this process is taken by many to provide support for the veracity of the model, or at least for a broader theory which the model exemplifies.

Despite the patina of rigor provided by the specificity and formality of the modeling, this should not be mistaken for genuine empirical testing (Roberts & Pashler, submitted). A close fit does not demonstrate that the model excludes any possible outcomes, much less any plausible ones. Even when there are observations the model could not have fit, the credit may belong not to the underlying theory but rather to the machinery added to allow "predictions" to be made (this machinery is not restricted to the so-called free parameters of the model, but often includes many aspects of a model that are not required by the basic theory). Third, even in those rare cases in which the theory really does predict the outcomes obtained, the enterprise is weakened by the failure to consider competing models and their predictions. Unless one has reason to believe that significant competing models do not make identical predictions, the model's victory may be chimeric. In cases where investigators have reexamined complex models backed up by data fitting, they often find that the data can be fit just as well using models that make radically different assumptions (e.g., Coltheart and Coltheart, 1972; J. C. Johnston et al., 1985). Evidently, then, data fitting should not be confused with theory testing.

The fact that the two are so often confused in scientific psychology is peculiar, especially because data fitting has fallen into poor repute in other fields. Unfortunately, it seems that psychologists sometimes mistake formalism and (apparent) precision in a theory for rigor in the empirical testing of that theory. These are two completely different things; whereas precision is a desirable feature in knowledge, it is not obviously a desirable feature in untested speculation. The data-fitting approach criticized here should not be confused with the field of mathematical psychology, which includes a greater number of contributions that derive nonparametric predictions from large classes of models, thereby facilitating genuine theory testing (see Townsend and Ashby, 1983, or Schweickert and Townsend, 1993, for examples relevant to attention).

In this book, the focus is on information-processing approaches to the phenomena associated with attention, undertaken without an assumption that the mind contains a structure or process that corresponds to that term. The emphasis is on experimental work that evaluates competing hypotheses, especially research that attempts to derive and test distinctive predictions of broad classes of models, because this is the most efficient way to home in on the truth. Complex models that have been fitted to data only in a post hoc fashion are not discussed. The aim is to see what functional implications can be drawn from the experimental literature on attention. Proposed architectures for attention that have been presented as purely speculative or have been justified only by data fitting are not addressed.

Background: Modern Theories of Attention

Now this chapter examines two important theories that emerged in the 1950s and 1960s and that framed a great deal of work that has taken place in the intervening years: early selection theories and late selection theories. This modern phase of attention research was inaugurated by post-World War II researchers in Great Britain, notably, Colin Cherry, Donald Broadbent, and Alan Welford. The major findings of these investigators are described in chapters 2 and 6. The transition from the early theorizing about attention to this modern phase was marked by a number of fundamental changes. One change, already described, was methodological: from theories based primarily on introspective observation to theories based on various kinds of behavioral data. A second change reflected the emergence of new concepts that permitted the events that make up human perception to be fractionated. Writers such as James and Titchener spoke of the conscious percepts elicited by stimuli and speculated on how attention might affect different attributes of the percepts, such as clarity or intensity. They did not, however, view perception as a process of achieving successively more elaborated and defined internal descriptions of a stimulus. In particular, they had little to say about the relation between attention and stimulus *identification,* an issue that emerged in the 1950s as central in modern attention theory (at a cost of neglecting other important issues; see Shulman, 1990). In fact, the idea of identification was not unknown in pretwentieth-century psychology; Hoffding (1896), for example, made a clear distinction between having a visual percept of an object and recognizing its identity. However, this distinction played little role in the most influential early analyses of attention.

The early and late selection theories that we are about to discuss dealt primarily with the perceptual aspects of attention, and the findings described in the next three chapters are critical for appraising these theories. In chapter 5 these theories are reexamined in light of that evidence. In addition to describing the contents of these accounts, the rest of this chapter discusses the logical relationship between different components of these theories. Those familiar with the history of attention research may believe that these theories have grown stale over the past few years, which is certainly a reasonable sentiment. It is nonetheless essential for readers to be aware of them in examining the evidence discussed in the first part of the book, for several reasons. One reason is that these theories motivated much of the research on attention, so the thrust of modern studies is hard to grasp without knowing something about them. The second reason is rather puzzling: although these accounts engendered an enormous volume of literature, the logical relationship between them and the broader space of possibilities of which they are a part has rarely been examined, and never to my knowledge in a systematic fashion. When one does examine it systematically, one can arrive at a more inclusive and coherent framework than the commonly recognized dichotomy (or occasionally trichotomy) of attention theories.

Early Selection Theories

The earliest and best known modern theory of attention was Broadbent's (1958) Filter theory, often referred to as early selection theory. Broadbent's basic hypothesis was that all stimuli reaching the sensory system are processed to the point at which certain physical attributes (e.g., location, loudness, and pitch of auditory stimuli) are analyzed and explicitly represented. He suggested that the machinery that identifies stimuli (e.g., recognizing an individual letter or word and comprehending its meaning) is capable of handling only one stimulus at a time. Thus, he postulated a filtering device responsible for determining which stimuli are to be processed further, and thought that it worked on the basis of the preliminary analysis of simple physical attributes. The mechanism responsible for selection he termed the selective filter—the workhorse of attentional selection. The term early selection theory and Broadbent's term filter theory are now commonly used to refer to this entire set of propositions. Early selection conveys the idea that all selection occurs early in the stream of processing. It is critical to note that the word early does not refer to time, but rather to the sequence of processing stages: selection is said to precede stimulus identification. Early selection theories are often represented with a structural diagram like that in figure 1.1 (Kahneman, 1973).

As the figure shows, all stimuli proceed far enough into the system to reach the first box (representing analysis in terms of physical features), whereas only the single selected stimulus proceeds farther, to the point

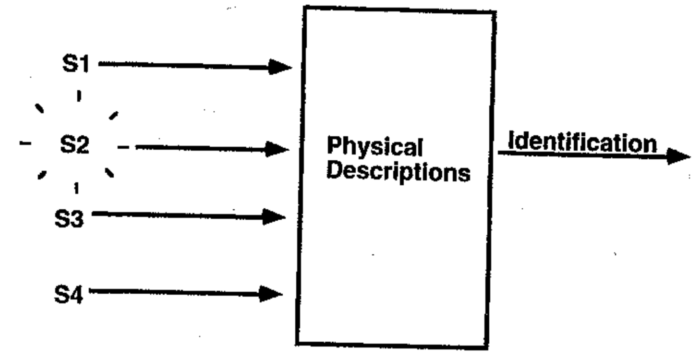

Figure 1.1
Early selection theory. The physical properties of the attended stimulus (S2) and the unattended stimuli (S1, S3, S4) are computed, but only the attended stimulus (S2) is identified.

of recognition. This kind of representation is not particularly apt for considering what early selection theory would imply for the time course of processing when a person attempts to perceive many stimuli. To make these implications more explicit, it is helpful to use a diagram in which time runs along the horizontal axis, as in figure 1.2. In this figure and subsequent ones, various stimuli presented to a person's sensory receptors are referred to with the abbreviations S1, S2, and so on. The physical attributes of the corresponding stimuli are referred to as PA1, PA2, and so on. Once a stimulus has been recognized, the internal representation that corresponds to its recognition may be depicted as an identity code, and the identity codes corresponding to S1, S2, and so on are abbreviated ID1, ID2, and so on. As an additional convention, highlights are placed around any stimuli the subject attempts to select for awareness and further processing.

The top panel in figure 1.2 shows the implications of early selection theory for the time course of processing two stimuli, of which only one, S1, is to be deliberately attended. According to the theory, the system works out the physical attributes of both S1 and S2 immediately, and then the identity of S1 (but not S2) is determined. The bottom panel shows the sequence of processing that must occur when two stimuli S1 and S2 are both attended; that is, the individual attempts to recognize

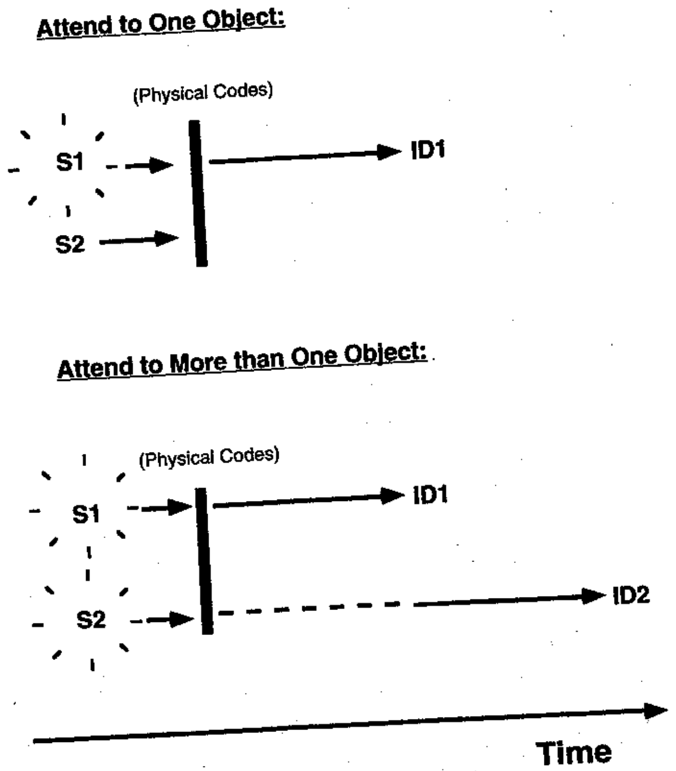

Figure 1.2
Time course of processing, according to early selection theory, when only one stimulus is attended (top panel) or two stimuli are both attended (bottom panel).

both of them. Perceptual processes, through the point of retrieving the physical attributes, operates in parallel; more elaborate processing involved in identifying stimuli occurs serially, with one stimulus analyzed and then the next.

Late Selection Theories

The most extreme alternative to Broadbent's early selection theory is what is usually called late selection theory. Well-known versions of the theory were proposed by Deutsch and Deutsch (1963), Norman (1968), MacKay (1973), and more recently by Duncan (1980b). These formula-

tions differ in interesting ways, but they have one basic idea in common: recognition of familiar objects proceeds unselectively and without any capacity limitations. One cannot voluntarily choose to identify or recognize something, according to these theorists. Whether there is just one sensory input or many does not affect the extent to which stimuli are analyzed or the timing of such analyses. Selective processing, which *is* subject to capacity limitations, is assumed to begin only after analysis is completed. Since by most accounts people seem to be unaware of many stimuli they attempt to ignore, such theories naturally suppose that awareness depends on these subsequent mechanisms. Duncan (1980b), for example, envisioned a transfer process that sends the results of an (unconscious) analysis process to further mechanisms. Enduring memory for what was perceived, and the ability to make a voluntary response are also assumed to depend on this transfer.

Figure 1.3 shows a typical way of depicting late selection theory. It is nothing more than a relabeling of the early selection diagram, with the label "physical descriptions" replaced by "semantic descriptions."

Several points about this theory—which may or may not be obvious—should be noted at the outset. First, it obviously does not claim that all the stimuli transmitting or reflecting energy to a person's sensory

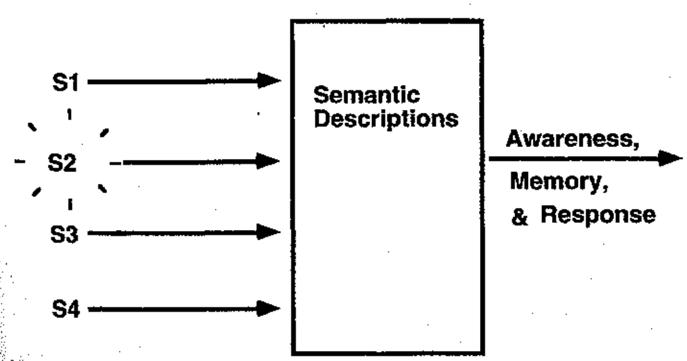

Figure 1.3
Late selection theory. The identity of both the attended stimulus (S2) and the unattended stimuli (S1, S3, S4) are computed alike. However, only the attended stimulus (S2) is selected for access to a system required for awareness, memory and response.

apparatus are identified. Limitations having nothing to do with processing capacity or attention may prevent this from happening. For example, a visual object in the periphery of the visual field may reflect a pattern of light to a part of the retina that contains too few receptors to pick up the information necessary for the object to be recognized, whether or not it is attended. This kind of recognition failure is not an attentional effect, and it does not contradict late selection theories. The same is true of *masking* (roughly, the harmful effects of one stimulus on the perception of another). For example, a visual stimulus that would otherwise be recognizable may fail to be identified because a stimulus follows it in the same part of the visual field (backward visual masking) or adjacent to it (lateral masking). In both cases, masking prevents a target stimulus from being identified even when the observer has every opportunity to attend to the target and tries hard to ignore the masking stimuli. Therefore, the assertion of late selection theory that stimuli are analyzed independent of attention or capacity limits should not be taken to mean that stimuli can always be identified; rather, it claims that voluntary control will have no effect on whether or not they are identified.

What does it mean to say a stimulus is fully analyzed? Modern attention theorists generally understood this to refer to whatever modes of categorizing a stimulus the person has practiced extensively. Categorizing letters, words, and speech into their respective linguistic categories, and categorizing objects such as tables and chairs in their basic-level categories are examples of what these theorists had in mind. It is doubtful that any late selection theorist intended the theory to cover whatever arbitrary descriptions a person *might* apply to a stimulus (the second-from-the-largest square, the only word on the page that contains more vowels than consonants, etc.). This point was implicit in most late selection formulations, however, and made explicit in only a few (e.g., Shiffrin and Schneider, 1977; LaBerge and Samuels, 1974).

Oddly enough, although late selection theorists talked mostly about linguistic descriptions of written and spoken language tokens, most of them had almost nothing to say about understanding language at the level of sentences and propositions. Many students, when they first hear about late selection theory, immediately point out that a person cannot understand two streams of speech at the same time. How then, they ask,

could anyone suppose that semantic analysis proceeds unselectively and without capacity limits? Do late selection theorists suppose that the meaning of the unattended sentence is computed unconsciously? The only one to address this issue explicitly was MacKay (1973). He proposed that although comprehending the meaning of individual spoken words was not subject to voluntary control or capacity limitations, understanding sentences requires storage in long-term memory, which, he suggested, is possible for only one sentence as a time. MacKay's suggestions are unique, however; otherwise, late selection theorists seem to have viewed single-word recognition as the paradigmatic case of semantic analysis, a usage that would strike a linguist, for example, as rather peculiar.

Late selection theory, then, amounts to the idea that irrespective of what a person might choose to attend to or ignore, the neural machinery that recognizes stimuli as belonging to familiar categories performs its computation for all incoming stimuli to the degree the sensory input is adequate to permit this. The implications of this theory for the time course of processing for attended and unattended stimuli are shown in figure 1.4. When S1 and S2 are presented at the same time, both are processed to point of being identified in parallel, whether only one is attended, as in the top panel, or both are, as in the bottom panel.

Major Alternatives

Readers who have not encountered these two classic theories of attention and the large literature that has sprung up around them will undoubtedly notice that they make fairly extreme claims. This is widely recognized, but many attention studies were and continue to be interpreted as favoring either early or late selection theory, as though it was somehow given that one of them had to be correct. Since the early 1970s a number of alternatives have been suggested, and most investigators seem to suspect that some form of compromise theory will be necessary. What is not widely agreed is how the space of plausible alternatives should be viewed. The goal of this section is to provide a rough taxonomy of the major possibilities. No attempt is made to be exhaustive or detailed; as in Twenty Questions, it is not efficient to spend a great deal of time elaborating on possibilities that are shortly to be ruled out by empirical

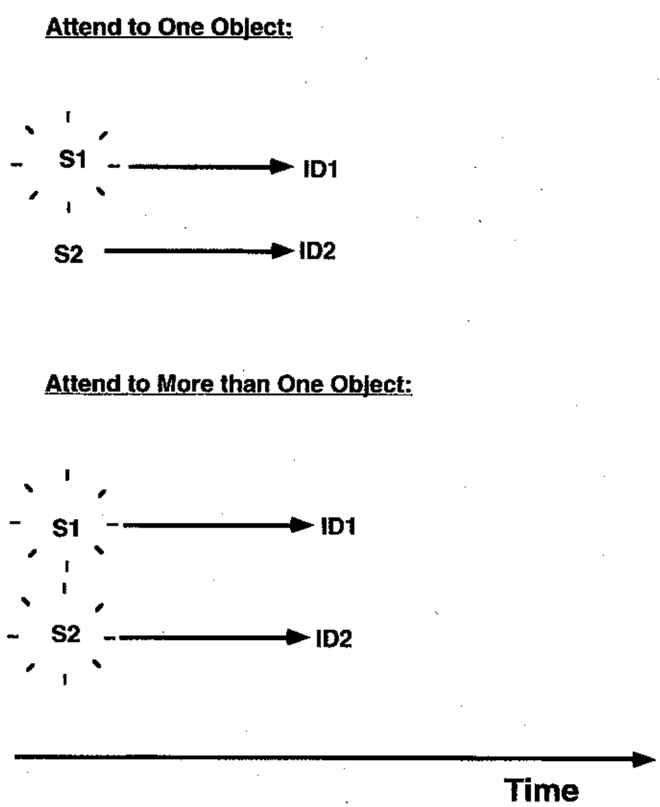

Figure 1.4
Time course of processing according to the late selection theory, when only one stimulus is attended (top panel) or when two stimuli are both attended (bottom panel).

evidence. However, it is worth considering what some of the possibilities would look like. The next section provides a rather breezy tour of the logical space of possible theories. When we return to general issues of attention theory in chapter 5, this space will be reexamined against the backdrop of empirical evidence.

One place to start is by noting a glaringly obvious alternative theory that has received surprisingly little explicit consideration. This possibility, often raised by undergraduates at their first introduction to the field of attention, is called the *controlled parallel* theory, and it is shown in figure

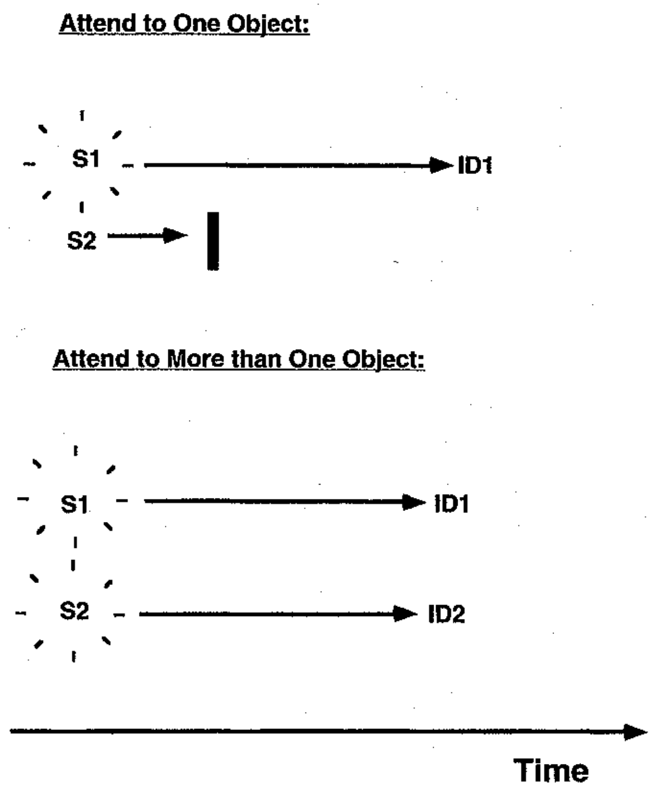

Figure 1.5
Controlled parallel processing, a natural alternative to early and late selection theories proposing parallel or serial identification depending on which is advantageous.

1.5. The basic idea is that when one stimulus is attended and another is rejected, the rejected stimulus is not analyzed beyond the physical level, as in early selection theory. When two stimuli are attended, both are identified in parallel, as in late selection theory. This possibility is an obvious one because it postulates that the system carries out the processing that would be most advantageous.

What is the relationship between the controlled parallel model and the two classic theories? All three can be seen as arrayed on a continuum, but it is not clear what single dimension of difference might be. It seems misleading, therefore, to describe the controlled parallel model as an intermediate theory. A more accurate view is that two separate questions are addressed by all three theories. The first question is, are rejected stimuli analyzed fully? The second is, can multiple *attended* stimuli be processed simultaneously when that is advantageous? Viewed in this way, all three theories are seen to reside within a two-by-two matrix. Figure 1.6 shows this matrix. The columns correspond to answers to the first question, and the rows of the matrix correspond to answers to the second. Late selection theory offers a positive answer to both questions, and early selection theory offers a negative answer to both. The controlled parallel theory offers the answers no and yes, respectively.

Having created a two-by-two matrix with three cells occupied, one naturally wonders why the lower-left hand cell is empty. To fit into this cell, a theory would have to claim that unattended stimuli are invariably identified (lack of selectivity), yet when two stimuli are attended, they must be identified one at a time. This requires the rather odd notion of a sequential process that keeps on marching—identifying stimuli one after another—even when there is no intention to analyze them. Is this account so strange that it can be dismissed? Perhaps not.

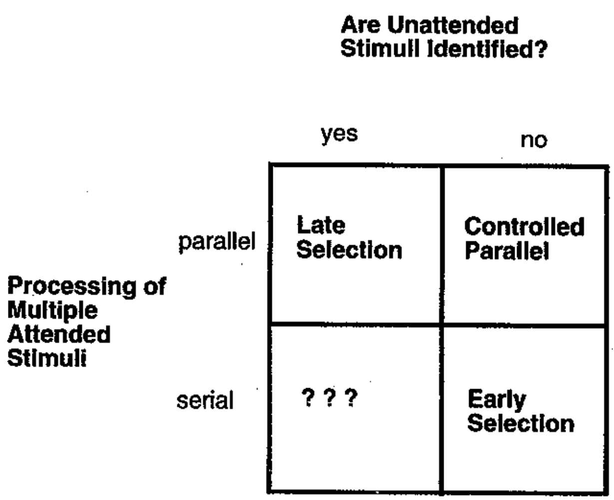

Figure 1.6
A two-by-two matrix of alternatives represents a simplified version of the space of possibilities.

When someone classifies a single stimulus and makes a rapid response that depends on the result, evidence about this *particular* stimulus is probably collected even after there is enough evidence for a response to be chosen (Rabbitt and Vyas, 1981; Levy and Pashler, 1995). For that reason, subjects are usually aware of the errors they have just made in speeded classification tasks (Rabbitt, 1979). Similarly, Saul Sternberg's well-known analysis of short-term memory scanning proposed that when people must decide whether a probe item matches any elements in a previously heard memory set, the comparison of the probe with items in memory continues even after a match has been found. This idea does not stand up well to empirical test (Monsell, 1978; McElree and Dosher, 1989), but few dismiss it as preposterous. Therefore, the idea that perceptual analysis might operate sequentially but exhaustively surely deserves at least a name. Henceforth, it will be called the uncontrolled serial model. What we are left with, then, is a two-by-two matrix with possible theories in all four cells, rather than the set of two extreme alternatives often envisioned in discussions of attention theory.

Elements of Compromise: Attenuation and Sharing

Of course this two-by-two matrix still sorts theories at a very coarse level, omitting from consideration many obvious intermediate possibilities. We turn now to some other theories that have been considered from time to time and demonstrate that their space cannot be captured fully in any matrix, no matter of how many dimensions.

One natural idea was suggested by Treisman (1960), who proposed that rejected messages are *attenuated* rather than completely blocked (the evidence for this idea is discussed in chapter 2). Treisman's idea had two parts. The first was that rejected stimuli are only filtered out partially, rather than completely. The second was an idea about when and how this partial information might be handled. Treisman suggested that recognition takes place through accumulation of information or activation in detector units. Unattended (and therefore attenuated) stimuli would not produce enough activity to cause the corresponding detector to reach its threshold. However, when the detector represents a concept that is somehow related to concepts that have recently been activated, which Treisman called *priming*, partial activation might suffice for recognition.

The implications of this filter-attenuation theory for *divided* attention were never clear. For example, would attending to more than one stimulus at the same time cause all the stimuli to be attenuated in the same way as an ignored stimulus would be? This question will be discussed further; for the moment, we can represent the idea of attenuation schematically as in figure 1.7.

A second widely recognized idea with elements of compromise between early and late selection theories is *graded capacity sharing*. This refers to the possibility that recognizing stimuli takes mental capacity or resources, and the total amount of this capacity is limited. Capacity might be shared among different perceptual processes, thereby reducing the amount available for any individual stimulus and causing its recognition to take longer. Figure 1.8 shows the theory as it applies to the case of one and two objects, *both attended*. To make it easy to think about capacity sharing (at the price of narrowing the concept further than one ought to), this figure follows McLeod (1977) in assuming that a given task requires a fixed total amount of capacity X time. This simplifying assumption allows one to represent the capacity allocated to a process at a given moment as the width of a channel. With time proceeding from left to right, the area of the channel represents the total capacity X time (i.e., work). If a process demands a fixed allocation of capacity X time before it can be completed, then sharing capacity equally between the two processes will double the time required for each, as shown in the figure. Obviously, this assumption is far too restrictive to be taken seriously, but it makes the general idea quite clear.

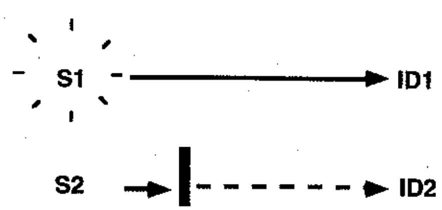

Figure 1.7
Filter-attenuation theory. According to this hypothesis, the information from the unattended stimulus is reduced in quality or intensity, although partial identification may still occur.

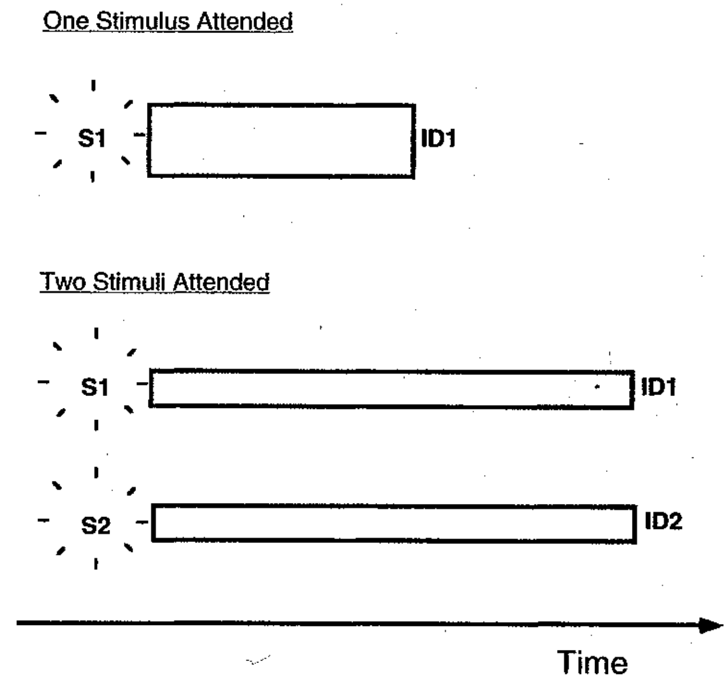

Figure 1.8
Capacity sharing. When two stimuli are attended, processing resources (represented by the height of box between stimulus and identity abbreviated ID) are divided, resulting in slower processing.

Various formulations of capacity-sharing accounts have been proposed. Kahneman (1973) suggested that allocation of finite resources might account for a broad range of limitations people have in doing different activities at the same time (in his version, the same resources were shared not only by perceptual activities, but also by cognitive and motor control processes). Accounts specifically focused on perceptual recognition and comparison were proposed by Townsend (1974), who showed that mean response times in visual search tasks could be accounted for in terms of capacity shared among different perceptual analyses. Shaw and Shaw (1977) and van der Heijden (1975) proposed somewhat similar formulations.

Proposals of this type, regarding possible capacity sharing in perceptual recognition, focused almost entirely on divided rather than selective

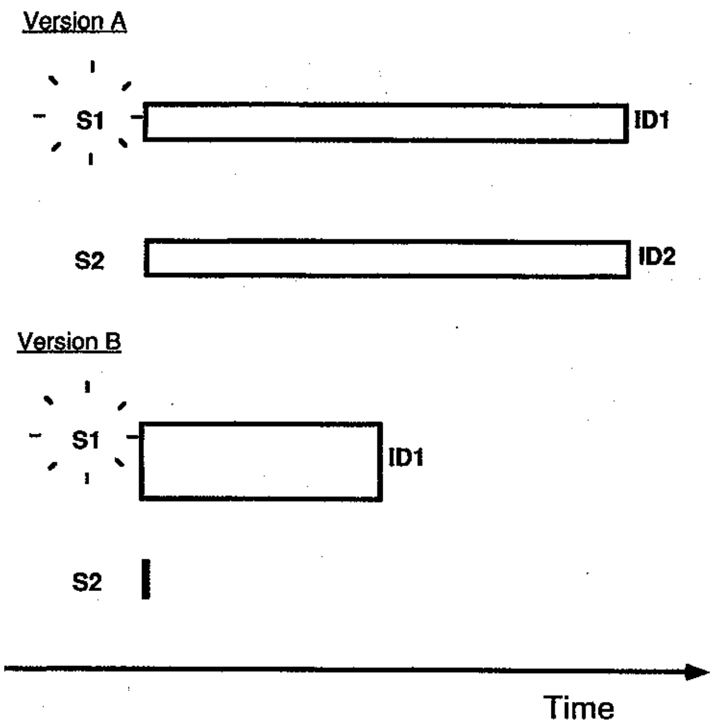

Figure 1.9
Two possible versions of capacity sharing when one input is attended (S1) and another is not (S2). In version A, the unattended item takes capacity; in version B, it does not.

attention. Thus, none of them had much to say about whether or not *rejected* stimuli also grab capacity, although it seems clear from Kahneman's (1973) formulation that he assumed they did not. In figure 1.9, Kahneman's version B depicts capacity divided exclusively among the attended items. The top panel shows the alternative version A, according to which the unattended stimuli also receive full capacity.

How does capacity sharing relate to early and late selection theory? The notion of capacity allocation is sufficiently general that other theories can be subsumed under it if one chooses. Early selection theory claims that processing capacity is allocated on an all-or-none basis to one stimulus at a time, and, one naturally assumes, generally only to the

relevant stimulus. Late selection theory maintains that there are enough processing resources for all the stimuli that could ever arise to be processed at maximum efficiency, and these resources are never denied to any stimulus. (Once one supposes that, however, there is not much gain in speaking of resources in the first place.) The general point is simply that capacity sharing is consistent with either early or late selection theory, but also allows intermediate possibilities.

Capacity sharing also points toward a whole class of intermediate theories that differ qualitatively from the classic ideas. Suppose limited resources can be simultaneously shared among different stimuli. The two panels in figure 1.9 represent two versions of such an intermediate theory; they differ only with respect to the locus of selectivity. The top panel is most naturally seen as a version of late selection theory modified to postulate graded capacity sharing, and the bottom panel is most naturally considered a version off the controlled parallel model, subject to the same modification.

At this point it is probably best to call a halt to the proliferation of theoretical alternatives. The possibilities discussed so far illustrate some of the most natural, or most influential, ideas about selection and capacity limits. With a little ingenuity, one could go on without limit, entertaining ever more baroque possibilities from one's armchair. Thus, one might combine capacity sharing with Treisman's hypothesis of attenuation for rejected stimuli by supposing that rejected stimuli tend to grab any spare capacity not already allocated. Or perhaps rejected stimuli grab a fixed amount of resources, leaving the balance to be allocated in a controlled way, and so on. However, it is undoubtedly more efficient to examine the empirical findings of studies of selectivity, capacity, and set, and return to these theories once some empirical constraints have been set out.

This rather breezy tour suggests a few general conclusions. First, it is not easy to provide an exhaustive taxonomy of functional analyses of attentional control and capacity limits in perception. Second, classic theories differ with regard to several key issues, and—most critically—these issues are at least partly orthogonal to each another. It is odd how rarely this orthogonality has been discussed, as it has important implications. Since the theories differ on a number of orthogonal issues, one should never ask whether the findings of a particular study favor late selection or early selection. Rather, one should ask how they bear on more specific hypotheses such as, rejected stimuli are always processed to the point of recognition or, certain types of stimuli cannot be recognized without capacity limitations. Before turning to the empirical findings, an overview is provided of some of the kinds of tasks that have been most commonly used in studies of attention and perception, and some terminology is introduced that is useful in referring to these tasks.

Methods and Terms: Laboratory Measures of Attention

Chapters 2 through 4 examine the basic experimental findings on attention and perception. A perceptual experiment just about always involves instructing the subject to perform some sort of task, even when the instructions do not inform the subject about the real purpose of the experiment. Attention researchers have used many variations on a rather modest number of different tasks, but they have a bewildering number of terms for these tasks. Fortunately, a simple taxonomy of basic tasks will cover most of the research designs. Insofar as possible, we will use the most common names, but since so many terms exist, the ones selected are probably used by a plurality, rather than a majority, of writers.

Filtering Tasks

In the filtering task, subjects are presented with more than one stimulus at the same time. Their job is to report something (the *reported attribute*) about just the subset of the presented stimuli that satisfy a certain *selection criterion*. The selection criterion is a physical attribute such as location, color, or loudness. The reported attribute is usually some category to which the stimulus belongs, one that is dependent on its symbolic identity. Clear-cut examples of instructions that define filtering tasks are "read only the red letter" and "read whichever word is being pointed at by the arrow." Selective shadowing (immediate repetition of spoken material) is also a filtering task. "Shadow the message presented to the left ear" defines a filtering task in which the selection criterion is ear of origin, or, more precisely, location as derived from interaural intensity differences, and the identity of the spoken words is the reported attribute. "Shadow just the man's voice" is another filtering task, with pitch the selection criterion. This oversimplifies the situation, because there is typically spectral overlap between a man's voice and a woman's.

The term "filtering task" obviously should not be taken to prejudice the issue of whether or not filter theory (Broadbent's early selection theory) is accurate. It is a fact that people can readily perform many kinds of filtering tasks. If late selection theory is correct, it would have to postulate mechanisms that allow people to select stimuli based on arbitrary physical attributes. It might even postulate that selection functions best when the selection criterion is physical (Duncan, 1981). The critical claim of late selection theories for filtering tasks is that the rejected stimuli are identified just the same as if they were attended (see figure 1.3). The controlled parallel and early selection theories would deny this claim, whereas Treisman's attenuation hypothesis would modify it.

Several variants of the filtering task have been widely investigated. In one such variant, the selection criterion is not specified well in advance, as in the cases described above. Rather, it may be specified shortly before the stimulus, simultaneous with it, or shortly after. The effects of the timing of precues can help in exploring the time course of the selection process itself (chapter 2), but they do not address either of the questions noted in figure 1.6. Many studies used *postcueing* (selection cues that followed the stimulus) to examine questions about sensory storage (i.e., information that outlasts the stimulus). Sperling (1960) was the first to observe that cues that follow a visual display by a short fraction of a second can be used to select information from a brief sensory memory system capable of holding onto much more information than a person can ever report. Studies of this sort addressed the questions of how much information is available and what kind of information it is. This literature is discussed briefly in chapter 3. These studies have some bearing on issues related to early and late selection, but that relationship is somewhat indirect. Chapter 2 deals with selective attention tasks, and almost all of the studies described in that chapter are filtering tasks.

Monitoring Tasks

The second category of tasks also involves numerous stimuli presented at the same time. However, the question asked of the subject depends on the categorical identity of more than one of the stimuli. In general, the word *monitoring* tends to be used with auditory stimuli or with visual displays that involve multiple frames,[3] and *search* is used more when the task involves a single visual display.

Monitoring and search tasks are commonly referred to as *divided-attention* tasks, reflecting the idea that one has to divide one's attention over several stimuli to carry out such tasks. This label is intuitively appropriate, but in line with the discussion above it will be avoided since it presupposes a claim that has to be tested, namely, that something called attention is divided in these tasks but not in selective-attention tasks.

Something else that has created considerable confusion is the use of the term late selection to describe tasks requiring selection on the basis of the identity of a stimulus. Once one adopts this way of speaking, one has to say that the simple fact that people can satisfactorily perform a monitoring task shows that late selection is possible. But of course this is misleading: the mere fact that people can do these tasks certainly does not support late selection theory, which states that these identities are computed unselectively, in parallel, and without capacity limitations (see figure 1.3). Therefore, the term late selection is reserved for the theory.

What is distinctive in the theory is the claim that selection *in the filtering tasks* involves selective processing that takes place only *after* all stimuli have been identified. The fact that, one way or another, people are capable of selecting a stimulus on the basis of its identity ("name the digit in the display") provides no special support for late-selection theory. Any viable early selection account would also have to account for it, which it could do by supposing that such a task would require sequential processing. According to the late selection or controlled parallel theory, monitoring can be accomplished in parallel without capacity limitations. Capacity-limited processing will obviously predict some reduction in efficiency with increasing numbers of attended items. The converse is not true: impairments with more attended stimuli do not logically imply capacity limits, a less than obvious point that is discussed in chapter 3.

Tasks Involving a Single Stimulus

Whereas many cognitive psychologists tend to regard both filtering and monitoring tasks as fairly austere, from the perspective of psychophysicists they are elaborate compared with the basic detection and discrimination designs. These classic psychophysical tasks do play an important

role in attention research (they appear frequently in chapters 2 and 4). Therefore, it is necessary to introduce some standard terminology for them as well.

In these tasks, there is only one stimulus and no selection criterion at all, at least, none that is explicit. In the yes/no detection task, the subject decides in each trial whether or not a stimulus was present. This kind of task is often analyzed with the help of signal-detection theory (Swets, 1964). The basic idea of signal detection theory is that every time a trial occurs and a stimulus is or is not presented, the subject's decision about whether or not to report a signal is based on the output of (certain channels in) the sensory system, which can be represented as a scalar quantity. This quantity is assumed to be subject to some random variability—random from the point of view of the experimenter. Classic signal detection theory furthermore assumes that the noise has a Gaussian distribution.

The subject is assumed to set a threshold and to respond "yes" in case the decision variable lies above this threshold. Variation in the probability that a target is present, or in the payoffs for different outcomes (e.g., correct detection, false alarm) is assumed to result in changes in the threshold. Variation in signal intensity is assumed to alter the strength of the signal (i.e., the difference between the signal-alone distribution and the signal-plus-noise distribution). This analysis has been applied not only to detection of stimuli at threshold, but also to search and monitoring tasks (e.g., Shaw, 1984; Palmer, Ames, and Lindsey, 1993)

Another popular task for exploring detection near threshold is the two-interval forced-choice procedure. Here, a signal is always presented in one of two successive temporal intervals, and the subject is required to indicate in which interval the signal occurred. If one assumes that this decision is made simply by comparing the magnitude of the sensory variable during the two intervals and choosing the interval in which it was higher, one can use percentage-correct accuracy as a measure of detection performance. Detection tasks with at most a single stimulus are discussed extensively in connection with divided attention (chapter 3) and perceptual set (chapter 4). This provides an abbreviated overview of the most basic experimental designs that recur throughout the book. Of course, individual experiments tend to involve variations or refinements of the basic designs.

Organization of the Book

The organization of the next four chapters of this book has already been hinted at. Chapters 2 through 4 deal with three major areas of study involving normal human subjects in the field of attention as it relates to perception. Chapter 2 deals with selective attention, tasks in which subjects are presented with some information to which they try to attend, and other information that they attempt to ignore. The main questions are what determines the effectiveness of selection and how extensively the ignored stimuli are processed. Chapter 3 discusses divided attention, situations in which an individual tries to process multiple stimuli presented at the same time. The primary question is how much perceptual analysis of simultaneous inputs can be achieved when such processing is advantageous. Examples of questions considered in this chapter include the following: if a person hears a different spoken word in each ear, can he or she recognize both at once? Can a person read a word and detect the occurrence of a tone at the same time? Chapter 4 examines the phenomena of perceptual set. The processing of a stimulus can sometimes be affected by a person's expectations about that stimulus. These expectations might concern some property of the stimulus other than the one that is reported. For example, many studies assessed whether a person is better able to identify a brief stimulus if he or she knows in advance where it will be presented in the visual field. These expectations might also pertain to the same attribute as what the observer reports. A question, then, would be whether a person can identify a letter more quickly and efficiently if he or she knows the letter will be drawn from a small set of known alternatives.

Chapter 5 returns to the main theoretical issues of attention as they relate to perception. It reviews the main empirical generalizations that emerged from the studies described in the previous three chapters, and asks how the phenomena of selective attention, divided attention, and set can be understood in a common framework. The role of attention as a theoretical construct or set of constructs reappears in this chapter.

Chapters 6 and 7 address problems of divided attention that go beyond perception and recognition of stimuli, examining the limitations that arise when people try to perform more than one task at the same time (where "task" refers to cognitive judgments and production of motor responses, rather than just recognition of sensory input). Chapter 6 deals with dual-task performance, and Chapter 7 with attention and memory. There are several reasons for discussing perceptual aspects of attention (chapters 2–5) separately from postperceptual (i.e., cognitive, memorial, and motor-related) aspects. While the boundaries between perception and postperceptual processes are not always clear, the neuropsychological evidence shows that brain damage frequently results in impairments confined to one or other of these domains (McCarthy and Warrington, 1990). Perceptual versus postperceptual is therefore an a priori reasonable candidate for a fundamental division in cognitive architecture. More critically, however, the findings of attention studies described here (especially in chapter 6) specifically argue for a distinction between perceptual attentional limitations and more central limitations involved in thought and the planning of action. Nonetheless, the connections between perceptual and postperceptual processing are numerous, and many remain to be carefully mapped out in future research. Chapter 8 addresses three broad issues that have been significant in attentional research, especially in recent years: automaticity, mental effort, and cognitive control, including research on task set.

The overall thrust of this book is bottom-up: experimental findings first, empirical generalizations that go beyond particular paradigms second, and glimmerings of general theory last. However, experiments without questions are mind numbing, and the more empirical sections of the book are focused around questions posed at a medium level of abstraction. The reader familiar with the attention field will find some of the conclusions reached here broadly in line with mainstream views, but others may appear more idiosyncratic and even disagreeable, especially with regard to the concept of processing capacity and automaticity. The hybrid view of attentional selectivity advocated in chapter 5 probably falls in the first category. By contrast, the conclusions of chapter 4, which suggest that noise reduction provides a better account of cueing effects in single-item displays than capacity allocation, challenge many widely held views. Chapter 6 argues that central attentional limitations often reflect a discrete processing bottleneck, harking back to the earliest studies of divided attention in the laboratories of Craik and Welford. Chapter 7 analyzes the role of attention in different memory systems. Recent research in this area has focused on the putative role of attention in distinguishing between implicit and explicit memory storage, but a very different view is presented here, focusing on the relationship between attentional mechanisms and short-term memory. Readers knowledgeable in the field will undoubtedly, therefore, find some points they disagree with; I hope they will also find the argumentation explicit and empirical enough that the disagreements will promote empirical progress rather than confusion.

I

Attention and Perception