4 Social Cognition

Louise Pendry

KEY CONCEPTS

accessibility
accountability
automatic process
categorization
continuum model of impression formation
controlled process
dissociation model
encoding
goal dependent
heuristic
inconsistency resolution
individuating information
outcome dependency
priming
rebound effect
retrieval
schema
stereotype
stereotype suppression

CHAPTER OUTLINE

This chapter introduces the topic of social cognition: the study of how we make sense of others and ourselves. It focuses especially on the distinction between social processes and judgements that are often rapid and automatic, such as categorization and stereotype activation, and those which may require more effort, deliberation and control (for example, stereotype suppression and individuated impression formation).

Introduction

What is social cognition?
What kinds of processes can social cognition research help to explain?

We inhabit a hectic social world. In any one day we can expect to deal with many other people. We may meet people for the first time, we may go out with old friends, we may find ourselves in a job interview trying to make a good impression on our prospective employer, queuing in a supermarket to pay for groceries, waiting for a train on a busy platform. Even for those of us professing to live ordinary lives, no two days are exactly alike. So, precisely how do we navigate this complex social life? What social information do we attend to, organize, remember and use? These are some of the questions that interest social cognition researchers, and providing answers to them strikes at the very heart of understanding human mental life.

As we go about our daily schedules, we are busy 'doing' social cognition for real. So, just what is it? Essentially, the study of social cognition promotes a deeper understanding of the mental processes that underlie human social behaviour (Fiske & Taylor, 1991). As Fiske (2004, p. 122) puts it, 'Social cognition analyzes the steps in people's train of thought about other people'. If we think about it for just a moment, it has implications for a very broad range of human social phenomena and domains. What is especially intriguing about social cognition is that it taps into the kinds of questions we find ourselves asking. Questions such as:

- Why did I assume that the man at the coffee machine in the boardroom was the company director when he was in fact the secretary?
- Why did I assume that Dr Alex James would be male/white?
- Why is it that I expected Albert to be elderly?

- Why did it surprise me to discover that Hilda, my elderly neighbour, had a passion for car maintenance?
- Why did I take the time to talk to my new black work colleague and find myself subsequently re-evaluating my initially stereotypic impression of her?

The aim of this chapter is to pass on to you a little of what we have learned thus far about some of the main theoretical issues in the field (for more detail see Bless, Fiedler & Strack, 2004; Fiske & Taylor, 1991; Kunda, 1999; Moskowitz, 2005). This chapter will give you a flavour of some of the more established theories in the field, and consider both the classic and more contemporary research that such theories have generated in their quest to understand better the workings of the social mind.

Although the field of social cognition is extremely broad and vibrant, few researchers would deny that one recurring, overarching theme is the distinction between social thinking that is *fast and furious* and social thinking that is more *measured and precise*. Since the 1970s, significant developments in theory and methodology have meant it is possible for us to now focus independently on these different types of thinking, that is, the influence of *unintentional* (i.e., unconscious) and *intentional* (i.e., conscious) processes in human thought and behaviour (Posner & Snyder, 1975). You might see this distinction encompassed in the term *dual-processing theories*. Or, to put it yet another way (and the way we will mainly refer to it hereafter), it refers to the contrast between ***automatic*** and ***controlled processes***.

automatic process a process that occurs without intention, effort or awareness and does not interfere with other concurrent cognitive processes

controlled process a process that is intentional, under the individual's volitional control, effortful and entailing conscious awareness

In this chapter, we focus primarily upon this distinction as it applies to *stereotyping* (e.g., Bargh, 1999; Devine, 1989). Do we process information about members of social groups carefully and rationally, or do we instead make rash judgements on the spur of the moment? Understanding when and why we engage in automatic or controlled processing can tell us a lot about how we view our social world.

To make this kind of distinction a little clearer, consider the following passage:

> Simon tried to put nationalities to faces, according to stereotype. The group of brawny, over-tanned and over-jewelled men and women who had ordered Bordeaux rather than local wine should be German – prosperous, large and loud. Any table giving off a cloud of cigarette smoke should be French, just as a table of non-smokers, with more water than wine being drunk, should be Americans. The English loaded butter onto their bread and ordered the heaviest desserts. The Swiss ate neatly and kept their elbows off the table, alternating sips of wine and sips of water like clockwork. (Mayle, 1993, p. 234)

Simon's observations may not strike us as particularly unusual. We may not endorse the national stereotypes conveyed in the above quotation, or at least question their accuracy, but somehow, we know exactly what he is talking about. Despite the perils of stereotypical thinking, it is something that we are apt to find irresistible (e.g., Brewer, 1988; Devine, 1989; Fiske & Neuberg, 1990).

Did Simon pause for even a moment to consider if his snap decisions were accurate? Did he stop to consider that in fact several of the French contingent were not actually smoking? Or that at least two of the English group were nibbling abstemiously on fruit salad (no cream)? No. The questions for social cognition researchers are: how and why did he respond in this manner? What processes led him to these conclusions? In social cognitive terms, the above illustration demonstrates several steps in a process that will be the focus of the next section. Simon has:

- *categorized* each of the persons in the restaurant in terms of nationality (grouped them into discrete sets – here, nationalities – based upon perceived shared characteristics);
- *activated* the *content* of these categories (what we term ***schemas***: our expectancies about people belonging to such groups);
- applied these schemas in such a way that *confirms the stereotypes* of the groups (looked only for information that is consistent with his expectancies).

schema a cognitive structure or mental representation comprising pre-digested information about objects or people from specific categories; our expectancies about objects or groups; what defines them

Much of this occurred fairly spontaneously. Simon did not stop to deliberate more carefully. As we shall see, it is a fundamental tenet of social cognition research that we often process information in precisely this way, only moving beyond the obvious stereotype if motivated and able to do so (e.g., Fiske & Neuberg, 1990). Stereotypes such as those outlined above have a functional role to play in facilitating person perception. They are, as Bodenhausen (1990) has noted, a kind of cognitive shortcut, a simplifying rule of thumb or ***heuristic*** that serves us well a lot of the time, but not always (for more on heuristics, see also Chapter 7, this volume).

heuristic a well-used, non-optimal rule of thumb used to arrive at a judgement that is effective in many but not all cases; stereotypes are often said to function as heuristics

In this chapter, we will consider some of the research that speaks to the automatic and controlled distinction: when, why and how do we engage in automatic versus controlled processing in person perception and stereotyping? By the end of this chapter you will be well equipped to provide preliminary answers to these questions.

JUMPING TO CONCLUSIONS: THE AUTOMATIC PILOT WITHIN

What makes a process automatic?
Are stereotypes activated automatically?
What part do categories and schemas play in the process of judging and understanding others?
How does stereotype activation impact upon behaviour?

Read the following passage:

> A father and his son were involved in a car accident in which the father was killed and the son was seriously injured. The father was pronounced dead at the scene of the accident and his body taken to a local morgue. The son was taken by ambulance to a nearby hospital and was immediately wheeled into an emergency operating room. A surgeon was called. Upon arrival, and seeing the patient, the attending surgeon exclaimed, 'Oh my God, it's my son!'
>
> Can you explain this?

So, how did you do? Many people find this question impossible to answer (based on lab class demonstrations over a number of years, often more than 40 per cent of students simply cannot do it). Moreover, they are apt to generate a wide range of convoluted explanations (for example, the 'father' who was killed is a Catholic priest and the term 'son' is therefore being used rather loosely) other than the most obvious one (the surgeon is the boy's mother). Why do people have so much trouble and why do they generate such complex rationalizations? Essentially, they find it hard to overcome the automatically activated ***stereotype*** (i.e., surgeons are generally men). As we shall see, this tendency to activate stereotypes automatically happens an awful lot. We will now pay some closer attention to why this occurs.

stereotype a cognitive structure that contains our knowledge, beliefs and expectancies about some human social group

What makes a process automatic?

For a process to be considered automatic, several criteria are deemed necessary (e.g., Posner & Snyder, 1975): the process needs to occur without *intention*, *effort* or *awareness* and is *not expected to interfere with other concurrent cognitive processes*. For those of us who have been behind the wheel of a car for a few years, the act of changing gear would possibly meet these criteria. A controlled process, on the other hand, is one that is: *intentional*, *under the individual's volitional control*, *effortful* and *entails conscious awareness*. To continue the driving analogy, deciding whether it is safe to overtake on a busy motorway should (one would hope!) fulfil these criteria. Let us now consider how this automatic versus controlled distinction contributes to our understanding of the process of stereotyping. We will start by looking at social categorization.

Plate 4.1 *Did you assume that these surgeons were men?*

The pervasive nature of social categorization

Almost every doctor who saw and examined me, labelled me a very interesting but also a hopeless case. Many told my mother very gently that I was mentally defective and would remain so . . . nothing could be done for me.

(Christy Brown, 1954/1990, p. 10)

Christy Brown suffered from cerebral palsy and was considered mentally disabled until one day he snatched a piece of chalk from his sister and wrote some words with it. He went on to astound and defy the medical profession by becoming a widely acclaimed

categorization the tendency to group objects (including people) into discrete groups, based upon shared characteristics common to them

author, whose autobiography was made into a successful film, *My Left Foot*. Branded from birth as 'retarded and useless', he fought to overcome prejudice and ignorance. The consequences of ***categorization*** can, then, sometimes be rather unpalatable. Once we assign others to particular social categories, associated stereotypic information can dominate our judgements to a worrying degree. Nonetheless, it does seem unlikely that we can view others in total isolation from their obvious physical and social categories. This is the view adopted by many theorists who work in the area (e.g., Brewer, 1988; Fiske & Neuberg, 1990; Macrae & Bodenhausen, 2000), and we will now turn our attention to why they endorse such a position.

Categorization refers to the tendency we have to group objects (including people) into discrete groups, based upon shared characteristics. There are object categories for furniture, takeaway food and musical instruments, but also social categories for women, refuse collectors, children, rock stars and so on. It is a fundamental premise of the social cognition approach that such categories serve a very useful function (e.g., Allport, 1954; Macrae, Milne & Bodenhausen, 1994). Why do you think this is?

Consider the following thought experiment. Imagine a far-off planet, Zygon, a place where perceptual mechanisms and inferential strategies have evolved in a decidedly unearthly manner. One day, an inhabitant from Zygon lands her spaceship somewhere on planet Earth and begins her journey into the unknown. She will doubtless be faced with many new objects and life forms that we human beings would, effortlessly, be able to sort into people, buildings, animals, trees and so on. Not so the Zygonian. Devoid of the cognitive know-how to parse this new and complex social landscape into something more intelligible, she would eventually experience information overload. There would simply be too many stimuli to process, at least in any meaningful manner.

This ability to separate our social world into discrete social categories is therefore a vital adaptation that ensures we don't find ourselves in a similarly daunting position. Without it, each person we met (or each object) would be unique and need be treated accordingly. Imagine how much time and effort that would take. Stated simply, *categorization favours simplification*, which in turn renders the world a more orderly, predictable and controllable place.

So, having established why categorization is so useful, let's look in more detail at the evidence that it is an *automatic* process. In 1989, Devine published an influential article in which she argued that (1) knowledge about racial stereotypes is culturally shared, even by people who do not endorse such stereotypes, and (2) activation of this knowledge (i.e., stereotype activation) is an automatic process. Recall that the criteria for a process to be automatic include that it is unconscious and does not require intention, attention or effort. If stereotype activation is truly automatic, this should mean that any time the appropriate cues are present (e.g., age, race or gender), stereotype activation should *invariably* result. So, how might this be tested empirically? Devine

PIONEER

Patricia Devine (b. 1959) spent her undergraduate years at the State University of New York, graduating in 1981, *summa cum laude*. This was followed by an MA (in 1983) and a PhD (in 1986) from Ohio State University. Devine's research centres around the intrapersonal and interpersonal challenges associated with prejudice in contemporary society. Her early work on the automatic and controlled components of stereotyping (1989) has been extremely influential in the field. Recent research concerns include the relation between explicit and implicit prejudice and the processes that regulate the use of stereotypes.

(1989, Study 2) used what is known as a ***priming*** paradigm (see Bargh & Pietromonaco, 1982). We need to dwell a moment on what priming is and why the priming paradigm is such a useful research tool to enable us to test Devine's hypothesis (and indeed, many other related research questions).

priming activating one stimulus (e.g., bird) facilitates the subsequent processing of another related stimulus (e.g., wing, feather)

When a construct is triggered in memory and made *temporarily accessible*, this is called priming and the stimulus that leads to this construct being triggered is called the prime (Moskowitz, 2005). In concrete terms, priming or activating one stimulus (e.g., bird) facilitates the subsequent processing of another related stimulus (e.g., wing, feather) via a process known as *spreading activation* (e.g., Neely, 1977). Once a construct is activated, associated concepts are also triggered and attain a state of heightened ***accessibility***, even if they were not directly primed initially. Such concepts therefore require some kind of cue to render them momentarily accessible. To use an analogy proposed by Higgins, Bargh and Lombardi (1985), these concepts are like a battery that is running low but can be recharged in certain circumstances (i.e., when the appropriate environmental trigger is present). It should be noted that other concepts, such as strongly held political beliefs, are often perpetually well charged (aided, for example, by repeated exposure to political arguments in the press, or political debates with like-minded friends). Being in a state of permanently high charge, they are routinely more accessible. These are termed *chronically accessible* concepts (for a detailed review see Moskowitz, 2005). Here, though, we focus primarily on how priming makes concepts temporarily accessible.

accessibility the extent to which information is easily located and retrieved

In one measure of accessibility, known as a *lexical decision task*, priming stimuli (e.g., words or pictures) are often presented on a computer, usually very quickly. Participants are then shown a letter string that may or may not be associated with the prime

(and may or may not be a real word), and asked to decide if it is a real word or a non-word by pressing a computer key. A priming effect is obtained when participants are shown to respond significantly faster to real words preceded by an associated prime (i.e., are quicker to respond to *wing* after being primed with *bird*). The advantage of priming paradigms is that they usually indicate uncontrolled automatic processing. Participants' subsequent reaction times are not prey to intentional self-presentational strategies (i.e., wanting to show themselves in a certain, often socially desirable, light), as might be the case, say, with paper and pencil measures of stereotyping.

Now that we are clearer about how priming paradigms work and why they are so suited to the study of automatic processes, let's return to Devine's work. In her experiment, primes related to a stereotype were presented outside of participants' conscious awareness. In order to do this, she presented the primes outside of participants' parafoveal field (i.e., out of their direct line of vision). The primes Devine used were terms related to the black stereotype (i.e., labels such as *blacks*, *niggers*, and physical or trait characteristics including *poor*, *lazy*). The participants had been pretested for prejudice level: half were high in prejudice towards black people, whereas half were low. This distinction forms an important part of Devine's experimental hypotheses, as we shall see later. Devine presented some participants with a high proportion (80 per cent) of ethnically associated words, and other participants with a much lower proportion (20 per cent).

Following the prime, in an ostensibly unrelated second experiment, participants read a brief scenario and were asked to form an impression of a target person who engaged in ambiguously hostile behaviours (after a paradigm originally developed by Srull & Wyer, 1979). Why hostile? Because pre-testing had indicated that hostility was a very strong feature of the black stereotype (see also Duncan, 1976). None of the words used in the priming phase, however, was directly related to hostility. This is important because it suggests that if the prime exerts the predicted effects upon interpretation of the ambiguous behaviour, it is due to automatic stereotype activation rather than simple priming of the hostile construct (but see Research close-up 4.3, below, for discussion of this point).

Let's consider what Devine predicted and found. Devine reported that those participants who received the high proportion of ethnic primes rated the target person in the story significantly more negatively (e.g., as more hostile and unfriendly) than did participants who received the low proportion of ethnic primes. Recall that Devine's view is that stereotypes are activated automatically. If this is so, then we should find that participants activate the black stereotype in the priming phase of the study (unconsciously) and go on to use it (without awareness) in the second part of the study (when forming an impression of the target). This should translate into higher ratings of the target as hostile, following a black prime. What about the differing levels of prejudice among participants? This was a very neat twist: if these results are found in both high- and low-prejudice individuals, it is stronger evidence still that stereotype knowledge is culturally shared and that activation is indeed automatic. If the priming effect can be demonstrated even among individuals who do not endorse the stereotype, this is pretty good evidence that it happened automatically. If low-prejudice participants could have found some way of controlling this rather undesired response, they surely would have done so, since it is clearly at odds with their beliefs. In fact, Devine found that participants' prior level of prejudice made little difference to how susceptible they were to the ethnic primes (but see Lepore & Brown, 1997, and the discussion in Research close-up 4.3).

This study is one of a number that have investigated the so-called automaticity of stereotype activation (see also Banaji & Hardin, 1996; Perdue & Gurtman, 1990). The results seem to provide quite compelling evidence. Moreover, during the 1990s research in this area blossomed and the literature is now replete with evidence of the seeming automaticity of stereotype activation (for recent reviews see Bargh, 1999; Devine & Monteith, 1999).

So should we conclude that the case for the automaticity of stereotype activation is established beyond question? Perhaps not just yet. The situation regarding automaticity is actually rather complex and researchers themselves are divided in terms of how it is best interpreted (see Bargh, 1999; Devine & Monteith, 1999). Moreover, recent research has provided some important qualifications to the debate, as we shall see later in this chapter. For now, we will note that stereotypes are often automatically activated. The question we now consider is this: once a stereotype has been activated, what can happen next?

Schemas: The next step in the process?

Several years ago a British national daily paper ran an advertising campaign on television. The advertisements featured a skinhead running at speed towards a businessman. Plates 4.2a and 4.2b show two shots in the sequence; what do you think happened next?

Most people, when asked, assumed the next shot showed the skinhead mugging the businessman. Turn to Plate 4.2c, p. 73, to see what actually happened. The newspaper used this example to illustrate its commitment to impartial reporting – the need to get the full picture. Here it serves a useful educational purpose: it potently depicts what can happen once a category has been activated. Why did people jump to this conclusion? The answer lies in the spontaneous ***encoding*** of the situation. People see the skinhead, readily activate the pertinent skinhead schema (e.g., anarchic, violent) and arrive at the mistaken conclusion that he is probably about to behave aggressively. Encoding refers to the way in which we translate what we see into a digestible format to be stored in the mind (Fiske & Taylor, 1991).

encoding the way in which we translate what we see into a digestible format to be stored in the mind

This example illustrates that whilst it may be a useful strategy to leap to the first obvious conclusions when perceiving others, it is not always a sound one. The behaviour was somewhat *ambiguous*: there are many reasons why a person may be running in the direction of another. The important point is that, in this case, the activated schema *biased the interpretation of the behaviour* in line with the skinhead stereotype.

This tendency has been demonstrated in a number of laboratory experiments. Duncan (1976) showed white students a video

(a)

(b)

Plates 4.2a and b *Two stills of a skinhead in the* Guardian *advertising campaign.*

Table 4.1 *When do we rely upon schemas? (from Fiske & Taylor, 1991)*

Role schemas may dominate over traits (role schemas more informative)
Subtype schemas (business woman) may be used more than superordinate ones (woman)
Information presented early on can cue schemas (primacy)
We use schemas that attract our attention (salience)
We use schemas that have previously been primed (accessibility)
We use schemas consistent with our current feelings (mood)
We use schemas relevant to controlling outcomes (power)

featuring a quarrel between two protagonists, culminating in one shoving the other. The race of the person performing the shove and being shoved was varied. Half the participants saw a white person push either a black or a white person, whereas the other half saw a black person doing the pushing of either a black or white person. Later, participants were asked to describe what they had seen. Irrespective of the race of the 'victim', when the protagonist was black, 73 per cent of participants said he had acted in an aggressive manner, compared to only 13 per cent when the protagonist was white. Thus, exactly the same behaviour was encoded differently depending upon (and in line with the stereotype of) the race of the aggressor. Such studies reveal how schemas can bias the interpretation given to social events. Let's now consider the schema topic in more detail.

Once we have activated a category stereotype, we bring into play the *knowledge* contained within these structures: our schemas or stereotypes (Brewer, 1988; Fiske & Neuberg, 1990). Schemas are – stated simply – packets of pre-digested information we hold in our heads about objects or people from specific categories: our expectancies about objects or groups. As an illustration, consider the kinds of information that come to mind when the category 'class swot' or 'teacher's pet' is activated (e.g., the studious pupil in the class who does nothing but work and is readily described as *boring*, *shy*, *introverted*, *socially unskilled*, *never goes out*, *bookish*, *unpopular*, *generally disliked*). Clearly, several different types of information may be discerned, including, for example, knowledge about 'class swots' (what they typically do and don't do) and value judgements about them (their likeability, popularity, etc.) However, a schema should not be misconstrued as a long list of separate, unrelated items and attributes. Rather, it is a cognitive structure within which attributes are organized and relations between them perceived. Thus, we might perceive a relationship between the fact that 'class swots' are socially unskilled and don't go out much, or perhaps between the observation that they are boring and not very well liked.

So, a schema contains many different kinds of knowledge about a particular category. Armed with this knowledge, the process of impression formation is greatly facilitated, because schemas affect how quickly we perceive, notice and interpret available information (Fiske & Taylor, 1991; Kunda, 1999). We are apt to rely upon schemas for a number of reasons. Table 4.1 summarizes some of the main ones (from Fiske & Taylor, 1991).

So far, we have seen how schemas can influence what we pay attention to and the way in which information is encoded. Allied to this, schemas play an important role in the process of what we subsequently remember about others.

Schemas and person memory

Schema theory in social cognition draws heavily upon associated work in cognitive psychology suggesting that schematic representations aid the organization, ***retrieval*** and recognition of material in memory (Bransford & Johnson, 1972). Once a schematic expectancy

retrieval the process of recovering information from memory once it has been encoded

Plate 4.2c *Final still of the skinhead in the* Guardian *campaign.*

(stereotype) is activated, research has shown that we are often better able to recall information that is consistent, as opposed to inconsistent, with this schema (for a review, see Stangor & McMillan, 1992). A study by Cohen (1981) demonstrates this well. Participants were shown a video of a woman interacting with her husband. Before viewing the video, half were told the woman was a librarian, half that she was a waitress. Some behaviours were consistent with the stereotype of librarians but inconsistent with the stereotype of waitresses. For example, she wore glasses and listened to classical music (librarian stereotype). Other information was instead consistent with the waitress stereotype, but inconsistent with the librarian stereotype (drinks beer, affectionate towards husband).

In a later recall task, participants showed better memory for information that was consistent with the stereotype expectancy. According to researchers, this tendency is explained by the fact that consistent information fits better with what we expect to be true about a person (Hamilton, Sherman & Ruvolo, 1990). Our prior knowledge structures (schemas) help us to tie several pieces of new information together and link them to existing beliefs. Macrae, Milne and Bodenhausen (1994), for example, showed that memory for a series of traits associated with the doctor stereotype (e.g., *caring*, *upstanding*) was significantly improved when participants were given a stereotypic label (i.e., doctor) than when it was not provided.

However, research has not always found preferential recall for stereotype-consistent information. Stereotype-inconsistent information can also predominate in our recollections of others. Nonetheless, the schema is still playing a part. For example, if we have a very strong expectation that a colleague, Jim, is a kind, cheery sort, we will be rather taken aback if he one day seems to be acting in an aloof and serious manner. That is not the Jim we know, and we will be apt to remember this out-of-character (schema-inconsistent) behaviour as a result.

Many laboratory studies testify to this tendency. Hastie and Kumar (1979) gave participants a list of behaviours performed by a person whom they were led to believe was intelligent. Some of

(a)

(b)

Plates 4.3a and b *What kind of music do you guess the stereotypical librarian and waitress listen to? Do they drink beer or wine?*

the behavioural descriptions fitted with this label (i.e., were stereotype consistent), some did not (i.e., were stereotype inconsistent) and some were simply neutral with respect to the stereotype. When later probed for their memory, participants recalled more inconsistent than consistent information. So, what is driving this effect? Certainly, it seems plausible that information that violates our expectancies will grab our attention, but there is probably more to it than that. Precisely because this information is so out of step with what we expect, we face something of a cognitive struggle to reconcile it with what we already believe to be true (our pre-existing schema).

Hastie and Kumar (1979) argued that information that does not correspond to a prior expectancy is harder to comprehend. As a result, it is processed more deeply and it is ultimately more memorable. Later research does suggest that the process of ***inconsistency resolution***, as it is called, is one that does demand attention. Pendry and Macrae (1999) asked participants to form an impression of a target. Half did so with no distractions, but the other half were required to do so at the same time as memorizing a long string of digits (known as a digit rehearsal task). This extra task meant that they were unable to devote all their attention to the impression formation task. Hence, participants' ability to recall more inconsistent information was significantly diminished when they formed an impression of the target under such cognitive load.

inconsistency resolution the way in which we reconcile inconsistent information with a pre-established schema

How do we reconcile these findings? Stangor and McMillan (1992) conclude their review by saying that inconsistent information will be preferentially recalled when participants are motivated to be accurate and attend to all presented information. Under these conditions, participants will make considerable effort to reconcile inconsistent information, and will be more likely to recall it. However, when participants are not motivated to be accurate, or else are preoccupied with a distracting concurrent mental task, they are probably more likely to recall consistent information.

Whilst the picture appears complex, for our present purposes we need really only note one important point. When it comes to person memory, be it for information that is consistent or inconsistent with a prior expectancy, and whether the expectancy is provided at the encoding (information presentation) or retrieval (memory) stage, *schemas definitely matter*.

Schema activation and behaviour

In the late 1990s, a number of articles appeared that demonstrated a very intriguing phenomenon: behavioural responses (e.g., walking slowly) can be automatically activated in response to an activated stereotype-relevant word (e.g., 'wrinkle'). (See Research close-up 4.1.)

This same basic effect was subsequently demonstrated using a range of category stereotypes and trait concepts, and several behavioural consequences (e.g., intelligence tasks, interpersonal behaviour, memory performance: for a review, see Dijksterhuis

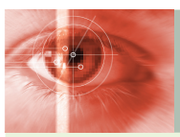

RESEARCH CLOSE-UP 4.1

The effect of priming on behaviour

Bargh, J.A., Chen, M. & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*, *71*, 230–244.

Introduction

This set of studies builds upon past priming research that demonstrates how a recently activated trait construct or stereotype, in an apparently unrelated context, can persist and exert an unintended effect upon the *interpretation of behaviour*. Here, the authors argue that *behavioural responses* to situations can also occur in response to an activated trait or stereotype prime. This rather disquieting suggestion is somewhat at odds with the prevailing assumption that behavioural responses to the social world are under conscious control. However, the authors reason that behavioural responses can be represented internally, just as are trait concepts and attitudes, and as such they should be capable of being automatically activated when triggering responses from the environment are present.

In three studies, the authors set out to put this hypothesis to the test. Here we focus on Study 2a: Behavioural effects of activating the elderly stereotype, in which participants were primed either with the elderly stereotype or with a neutral prime, and their subsequent walking speed was assessed. The authors hypothesized that elderly-primed participants would demonstrate significantly slower walking speeds in comparison to neutral-primed participants.

Method

Participants

Thirty male and female students participated in the study.

Design and procedure

Participants were first asked to work on a scrambled sentence task under the guise of a language proficiency experiment. For each of 30 items, participants had to use the five words listed to construct a grammatically correct four-word sentence as quickly as possible. This task formed the priming phase, serving to activate (or not) the appropriate stereotype. Hidden within

the scrambled sentences were words either relevant to the elderly stereotype (e.g., *grey*, *bingo*, *wrinkle*) or neutral, non-age-specific words. Participants were randomly assigned to either the elderly or neutral prime condition. Importantly, elderly words associated with slowness (a common elderly stereotypic trait) were excluded from the elderly prime condition. After completing the task, participants were partially debriefed. A second experimenter then covertly recorded the amount of time participants took to walk down the corridor after leaving the laboratory. Finally, participants were fully debriefed.

Results

After the conclusion of the experiment, participants in the elderly priming condition walked down a hallway more slowly than neutral prime control participants (see Figure 4.1).

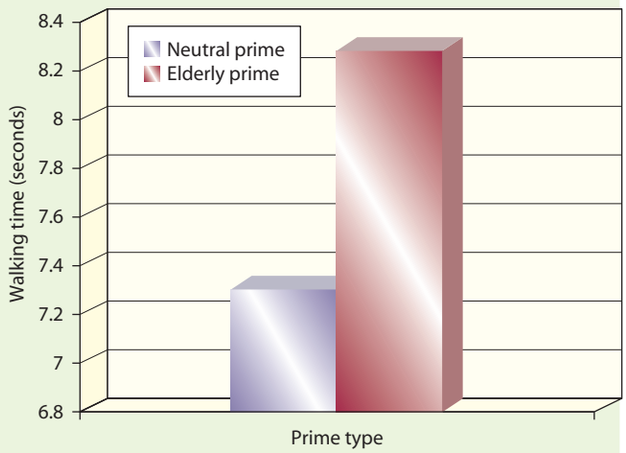

Figure 4.1 *Mean time (in seconds) needed to walk down a hallway as a function of prime type (from Bargh et al., 1996, Study 2a).*

Discussion

These results, together with data from other studies reported in this article, provided compelling initial evidence in support of the authors' hypothesis. After participants were exposed to an elderly prime, they demonstrated motor behaviour in line with the activated stereotype (i.e., slower walking speeds). Importantly, the authors took care to exclude any references to time or speed in the stimulus materials, so the effect is not simply a result of trait priming. This suggests that the elderly-prime stimulus words instead activated the elderly stereotype in memory.

The take-home message is that social behaviour can be triggered automatically by relevant features of the stimulus environment and can occur without awareness. This finding is qualified somewhat by the authors' observation that it may only occur if the behavioural representation activated by the prime is already associated with that situation, and that the motivation to act thus is part of the person's behavioural repertoire. As they note (Bargh et al., 1996, p. 240), 'It is doubtful, for example, that the participants in Experiment 2 left our building to go buy condos in Florida' (Florida being a popular spot for retired persons in the USA).

This paper spawned a great deal of interest and subsequent research. As an initial demonstration of what is termed 'automatic social behaviour', it remains a classic.

PIONEER

John Bargh (b. 1955) attended the University of Illinois as an undergraduate. He attended graduate school at the University of Michigan, with Robert Zajonc as his adviser. In 1981 he was awarded his PhD, and he then worked for several years at New York University. Currently he is at Yale University. He has received many prestigious awards, including the SESP Dissertation Award in 1982 and the APA Early Career Award for contributions to psychology in 1989. His research interests centre on the topic of automaticity and analysis of the unconscious nature of a wide range of psychological phenomena such as attitudes, emotions, motivations and social behaviour.

& Bargh, 2001). So, why does it happen? Again, automatic stereotype activation is thought to play a part. Think back to the work we reviewed earlier. Recall that once a social category (e.g., black) is activated, associated stereotypic traits (e.g., musical, hostile) are also activated (e.g., Devine, 1989). Likewise, trait words strongly linked to a particular category can act to cue the activation of the category. In the Bargh et al. (1996) study, for example, the trait primes associated with the elderly appeared to cue the activation of the elderly stereotype and lead participants to walk away from the experiment more slowly than non-primed participants.

Other studies have shown that category primes (e.g., professor) can cue the activation of associated traits and behaviours (e.g., intelligent, hard-working), leading to superior performance on an intelligence-related quiz (Dijksterhuis & van Knippenberg, 1998). So, when participants are primed with 'professor', this seemingly activates traits and behaviours associated with the category. Participants are therefore assumed to approach the subsequent quiz task with these traits unwittingly at the forefront of their minds, and the one that is relevant here (intelligence) exerts an influence upon performance.

Plate 4.4 *Primes associated with the elderly appear to cue the activation of the elderly stereotype (e.g., walking more slowly).*

These findings may at first seem somewhat implausible, but there is quite a lot of evidence accumulating in support of these results. In a recent review article, Dijksterhuis and Bargh (2001) provide a possible process explanation for their findings (see Figure 4.2).

They argue that the effects of stereotype activation on memory are mediated by trait activation. In other words, these effects are only shown when participants do activate the traits implied by the stereotype. Dijksterhuis, Aarts, Bargh and van Knippenberg (2000) showed that activating the elderly stereotype made participants more forgetful, but only if they actually associated the elderly with forgetfulness in the first place. Dijksterhuis and Bargh (2001) further note that trait concepts can activate behaviour representations. For example, activating the trait 'slow' results in activation of concrete behaviours such as 'linger' or 'dawdle'.

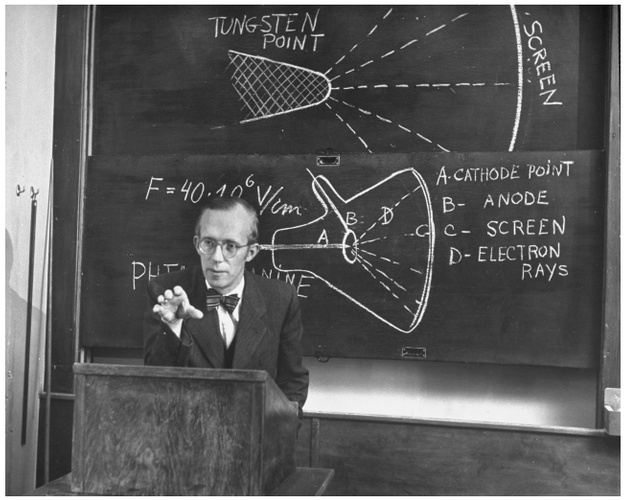

Plate 4.5 *Could being primed with a photo of a 'professor' make you perform better at a quiz?*

Research of this kind provides more evidence for the idea that stereotypes can be automatically activated and, in this case, lead to some rather unexpected consequences. Later on in this chapter, though, we shall see that such effects are not inevitable.

SUMMARY

Thus far, we have considered how stereotypes can be activated automatically and explored some of the consequences of this. Taken to an extreme conclusion, such research can generate some seemingly pessimistic conclusions about how much in control we are of our person-perception faculties. If so much goes on without our awareness, are we forever at the mercy of our processing frailties? That is certainly one interpretation of this literature (and for a spirited and compelling defence of this view, see Bargh, 1999). Many researchers, though, would argue that just because social information can be automatically activated, it does not necessarily follow that we will act in line with this information (Devine, 1989; Fiske, 2004; Monteith, Sherman & Devine, 1998; Monteith, Spicer & Tooman, 1998).

So, what factors cause us to look beyond our first schema-driven impressions and instead engage in a more systematic appraisal of the data? In the next section we shall consider a number of interlinked theoretical approaches that will provide some answers to this question.

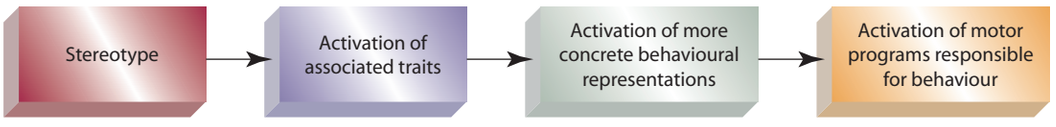

Figure 4.2 *A possible process explanation for the automatic social behaviour effect (from Dijksterhuis & Bargh, 2001).*

GOING THE EXTRA MILE: REGAINING COGNITIVE CONTROL

When are stereotypes not automatically activated?
Under what conditions does stereotype activation not lead to a stereotypic response?

Recently a colleague recounted the following story. He had been introduced at a party to a staff member from another department, computer science. He recounted his tale thus: 'When I first saw him I thought, computer nerd! The geeky 70s clothes, the old-fashioned specs, the terrible hairstyle. The kind of guy who drives a Skoda and watches endless episodes of *Star Trek*. But after a few minutes of talking to him, I had to think again. This guy was wild. A real party animal, who drove an Audi TT and enjoyed bungee jumping in his spare time. I couldn't have been more wrong.' This illustration hints at the yawning gap that can exist between our initial category-driven impressions and the reality of what lies beyond. The colleague took the time to get beyond his initial impression. In this section we'll look at when and why, like this colleague, we may go this extra mile.

Stereotype? What stereotype? Stopping the stereotype being activated in the first place

In the wake of early research implying the inevitability of stereotype activation, researchers have documented a number of qualifications to this view. Let's consider two indicative illustrations. Macrae and colleagues (1997) showed that stereotypes are not inevitably activated. Rather, perceivers have to process target information in a semantic manner in order to activate the stereotype (see Research close-up 4.2).

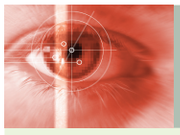

RESEARCH CLOSE-UP 4.2

The goal-dependent nature of stereotype activation

Macrae, C.N., Bodenhausen, G.V., Milne, A.B., Thorn, T.M.J. & Castelli, L. (1997). On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology*, *33*, 471–489.

Introduction

This article challenged the classic view that mere exposure to a member of a stereotyped group is sufficient to activate the associated stereotype. The authors reason that the degree to which stereotypes may be automatically activated may relate to the extent to which we are interested in the social meaning of the stimuli we encounter. It is possible that in some situations (for example, when you are in a busy lecture hall trying to locate your friend) you are more concerned with navigating your way through the throng than with forming impressions of the individuals who comprise it. In such situations, where people are simply objects to be navigated around en route to your desired destination, perhaps you do not activate any stereotypes at all. This situation was one that was modelled in a study by Macrae et al. (1997, Study 1).

Method

Participants
Forty-eight students took part in the study.

Design and procedure
The study had a 3 (processing set: feature detection or semantic judgement or exposure) × 2 (trait type: stereotypic or counterstereotypic) mixed-design with repeated measures on the second factor. Participants were shown faces of female undergraduates and pictures of common household objects. One group of participants was asked to detect, by means of a key press, whether a white dot appeared on each picture (feature detection condition). A second group was told just to press a key once each picture appeared on screen (exposure condition). A third group (semantic condition) was directed to process the pictures in a semantic manner: to decide whether each picture was of an animate or inanimate object (this is a semantic task because, in order to reach this decision, participants need first to process the object in terms of what it is and what it is called – this is semantic processing).

Participants also completed a lexical decision task (LDT). Each time they responded to a picture, a letter string would appear on screen and participants had to decide if it represented a real word (e.g., emotional) or a non-word (e.g., ingrac). The logic is that participants are quicker to respond to words if the construct associated with the words has previously been activated. So, if the construct 'woman' is activated in the first phase of the study, participants should be quicker to respond to words associated with this category (e.g., emotional) than if the construct has not been activated.

Results

The dependent measures of interest were latencies of response on the LDT task (see Figure 4.3). The results suggested that on trials where a photo of a woman had been presented, responses to stereotypically female word strings were much faster than to counterstereotypical words. It seemed that the picture of a woman had indeed activated the female category. However, this effect was only found for participants who had been instructed to process the photographs in a semantic fashion. So the simple feature detection goal (i.e., check stimuli for spots) served to eliminate stereotype activation. For activation to occur, then, some basic level of interest in the target had to be present.

Discussion

This study provides evidence that stereotype activation is not always a spontaneous by-product of a triggering stimulus. Rather, activation only occurred when participants processed the target in a semantic manner. The authors assert that the activation (and indeed, application) of stereotypes is likely governed by pragmatic concerns, here related to the particular processing goals in place. In sum, and in line with a growing research literature, this study highlights the goal-dependent nature of stereotype activation.

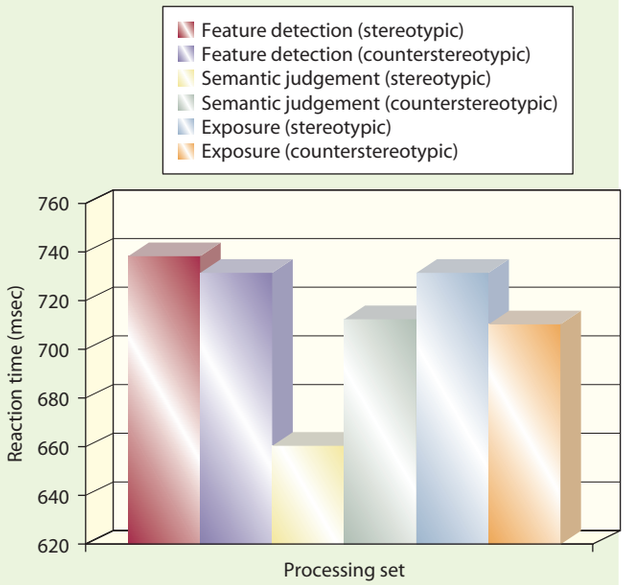

Figure 4.3 *Mean LDT (lexical decision task) latencies in msecs as a function of processing set and trait stereotypicality (from Macrae et al., 1997, Study 1).*
Note: Latencies for pictures of women (not objects) shown.

Stereotype activation is also affected by the extent to which participants endorse egalitarian world views. Moskowitz, Gollwitzer, Wasel and Schaal (1999) suggest there may be an effortless, preconscious form of cognitive control that in certain individuals prevents stereotype activation. Consider two people, Jack and Joe. Both would say they are low-prejudice, but whereas Jack would experience a feeling of incompleteness and self-disappointment upon learning that he had inadvertently acted in a stereotypic manner (and may want to do something about it), for Joe this realization would not be too troubling. Moskowitz et al. (1999) would view Jack as a 'chronic egalitarian': committed to being egalitarian, fair, tolerant and open-minded. Joe, on the other hand, would be more of a 'non-chronic egalitarian' in this respect. Do such differences impact upon stereotype activation? Moskowitz et al. (1999, Study 3) investigated this very question.

Participants were classified as chronic or non-chronic based upon responses to measures designed to assess commitment to egalitarian goals (here, with respect to fair treatment for women). In a second phase, participants saw photographs of men or women followed by an attribute, and were asked to pronounce this attribute as fast as possible. The attributes were either consistent with or irrelevant to the stereotype of women, and they were presented either 200 ms or 1,500 ms after the prime. Stereotype activation was demonstrated if participants were quicker to respond to stereotype-relevant attributes (e.g., kind) following stereotype-relevant primes (woman). Importantly, only non-chronics showed evidence of such stereotype activation. Participants with chronic goals failed to show this effect. This lack of activation could not, however, be due to conscious goals exerted on the part of chronics, because the difference between chronics and non-chronics was found even when attributes were presented 200 ms after the prime (conscious control is possible only after 600 ms have elapsed between the presentation of a prime and a stimulus). This led Moskowitz et al. (1999) to conclude that stereotype activation is not inevitable.

These and several other studies temper the initially extreme conclusions reached about the inevitability and automaticity of stereotype activation. In a review article, Macrae and Bodenhausen (2000) argued that two factors seem to play a prominent role in the regulation of category activation: perceivers' temporary processing goals and their general attitudes (i.e., prejudice level). The research outlined above is in accord with this view. Hence, category activation would appear to be ***goal dependent*** (Bargh, 1994), arising from the interplay of cognitive and motivational factors. In sum, evidence is accumulating that suggests it is sometimes possible to prevent stereotype activation (for more detail see Bargh, 1999; Devine & Monteith, 1999; Moskowitz, 2005; see Research close-up 4.3, p. 82).

goal dependent where an outcome is conditional upon a specific goal being in place (e.g., goal-dependent automatic stereotype activation)

Quashing the effects of stereotype activation once it has occurred

What happens, though, in those situations in which it is not possible to prevent activation? If stereotype activation happens, and

we would really rather it hadn't, what can we do? Most researchers agree that perceivers are able to exercise some degree of choice in their responses towards others (Fiske, 1989), provided they are *aware* of the potential influence of the stereotype, have sufficient *cognitive resources* available to exert control, and are in some way *motivated* not to respond in a stereotypic fashion (Devine & Monteith, 1999). If we fulfil these important criteria, then there are a number of strategies at our disposal.

In this section we consider several theoretical approaches that have furthered our understanding of the processes that may intervene following stereotype activation: Fiske and Neuberg's (1990) continuum model of impression formation; Devine's (1989) dissociation model of stereotyping; research on stereotype suppression (e.g., Macrae, Bodenhausen, Milne & Jetten, 1994); and moderators of the perception–behaviour link (Dijksterhuis & Bargh, 2001).

Impression formation: The rocky road from category-based to individuated processing When forming impressions of others, we commonly rely upon two sources of information: (1) knowledge of a person's category membership (e.g., female, elderly, as we saw in the work reviewed earlier) and (2) details of his or her personal or ***individuating*** characteristics (e.g., honest, forgetful). The persistent problem facing researchers has been to determine which of these contrasting sources of information contribute to the impressions derived (e.g., Brewer, 1988; Fiske & Neuberg, 1990).

individuating information information about a person's personal characteristics (not normally derived from a particular category membership)

continuum model of impression formation a theoretical model advanced by Fiske and Neuberg (1990) that views impression formation as a process going from category-based evaluations at one end of the continuum to individuated responses at the other. Progress along the continuum is thought to depend upon the interplay of motivational and attentional factors

Fiske and Neuberg's (1990) ***continuum model of impression formation*** provides one detailed answer to this puzzle. This model proposes that perceivers' evaluations of others fall somewhere along a continuum of impression formation, with category-based evaluations anchoring one end of the continuum and individuated responses the other. Constructed upon a number of theoretical premises, the model asserts that: (1) category-based responses have priority, and (2) movement along the continuum, from category-based to individuated responses, is a function of interpretational, motivational and attentional factors.

According to the model (and in line with much of the work we have already considered in this chapter), perceivers initially encounter a target and readily categorize him as a member of a particular social group. They then consider the personal relevance of the categorized target in the context of currently active concerns and goals. If the target is of little interest (e.g., the perceiver is merely passing a person in a street), then the impression formation process is short-circuited and resulting evaluations are predominantly category-based. If, however, the target is of at least minimal relevance (e.g., the target is an interviewer and the perceiver an interviewee hopeful of securing a new job), attentional resources are allocated to an appraisal of his or her personal attributes, and the protracted journey towards a more individuated impression begins. There are several stages at which processing can stop. An illustration of how this might work in practice in different situations is provided in Figure 4.4.

PIONEER

Susan Fiske (b. 1952) obtained her PhD from Harvard in 1978. After a number of years at the University of Massachussetts (Amherst) she moved to Princeton. In the course of her career she has amassed many prestigious awards, including (with Shelley Taylor) the 2003 Thomas Ostrom Award from the Person Memory Interest Group for work in social cognition. A past president of the American Psychological Society (2002–2003), she has published numerous articles, book chapters and books. Her current research focuses upon how stereotyping, prejudice and discrimination are encouraged or discouraged by social relationships, such as cooperation, competition and power.

Thus, initial categorization is relatively spontaneous, but the social perceiver will only stop here if the motivation to go further is lacking or if there are pressures (e.g., scarcity of time) conspiring against a more systematic appraisal of the evidence. Research on perceiver motivation and its effects on the impression formation process has resulted in the identification of several goals and task objectives that reliably elicit more individuated processing. Among the most important are: (1) ***outcome dependency*** on a target (participants believe they will later meet the target and work together on a jointly judged task; Neuberg & Fiske, 1987; Pendry & Macrae, 1994); (2) perceiver ***accountability*** (perceivers believe they will have to justify their responses to a third party and be held responsible for their impressions; Pendry, 1998; Tetlock, 1983); and (3) accuracy-set instructions (perceivers are instructed to be as accurate as possible; Kruglanski & Freund, 1983).

outcome dependency a motivational objective in which participants believe they will later meet a target and work together on a jointly judged task; shown to lead to less stereotypical target impressions

accountability a processing goal whereby perceivers believe they will have to justify their responses to a third party and be held responsible for their impressions; this typically leads to less stereotypical impressions

While differing on a number of counts, these motivational factors all share a common feature: they increase perceiver involvement with the target and encourage more individuated impressions. However, motivation to engage in controlled processing may on its own be insufficient if cognitive resources are depleted. For example, Pendry and Macrae (1994, Study 1) led participants to believe they would meet and interact with Hilda, an elderly female. Half the participants were also made outcome dependent: they stood to gain £20 for their joint performance with Hilda on a word-puzzle task. The remaining participants would work with

	Example	Example	Example
Initial categorization	Woman encountered in busy supermarket with a clutch of kids trailing behind: probably a mother *Processing stops here*, target is of no further interest/perceiver is in a hurry	Person overheard in next office, has high-pitched voice: probably female Perceiver needs to know more (could be new work colleague), looks for more clues	Elderly person who is female called Hilda. Perceiver needs to know more (it's his prospective mother-in-law)
Confirmatory categorization		Person is applying make-up: definitely female Perceiver still not satisfied, processing continues	Hilda enjoys listening to *The Darkness* and visiting her grandchildren: not your average elderly woman; hard to confirm initial categorization as sufficient, perceiver carries on
Recategorization		Wait a minute, she is also carrying a briefcase and a palmtop organizer, so actually she is probably a *business* woman *Processing stops here*, perceiver is satisfied (realizes he will not be working with this business woman); recategorization will suffice	Hilda services her own car and likes spicy curries and flower arranging: defies an obvious recategorization, perceiver probes deeper
Piecemeal integration			This family-loving, elderly woman called Hilda enjoys loud modern music, calmer creative activities, is mechanically minded and thrives on a diet of chicken vindaloo and Bombay potatoes *Processing stops here*: target is not amenable to a categorical impression; perceiver satisfied with impression, although he has an extremely atypical mother-in-law

Figure 4.4 *Fiske & Neuberg's (1990) continuum model of impression formation: an illustration of how processing can stop at different stages.*

Hilda, but their outcome would not depend on her performance. All participants received the same information about Hilda, half of which (12 items) was stereotypic, half of which was counterstereotypic.

Whilst reading the profile, half the participants performed a resource-depleting concurrent mental task (digit rehearsal), the others simply read the profile. To assess their impressions, all participants were asked to rate six personality traits (three pretested as stereotypic and three as counterstereotypic with respect to elderly females) for how characteristic they were of Hilda. Pendry and Macrae predicted and found that the formation of an individuated impression was contingent upon participants being both motivated (here, by being outcome dependent) and having full processing capacity (i.e., not being required to rehearse the digit whilst forming the impression).

In a second study, Pendry and Macrae (1994) sought to establish whether participants who are outcome dependent rather than outcome independent devote a greater proportion of their attentional capacity when forming an impression of a target. The idea that motivated perceivers allocate more attention to processing information is a fundamental premise of Fiske and Neuberg's model, although support for it at that time was somewhat limited.

To test this hypothesis, Pendry and Macrae (1994, Study 2) used what is called a *probe reaction task* (PRT; see Bargh, 1982). Participants were instructed to optimize their performance on the impression formation task and to use their remaining attentional capacity to respond to a subsidiary probe stimulus (i.e., turning off a randomly illuminated light bulb that appeared several times on a computer screen whilst the impression task was being performed). Importantly, this probe reaction task was not a method of resource depletion (like the digit rehearsal task). That is, its purpose was not to divert attentional resources away from the primary impression formation task and make the process harder. Rather, it assessed what attention was not being used in the primary task (how much attention was left over). If more involving motivational goals do entail greater attention to the target, then we should expect that participants under these conditions would have less attention left over to switch off the light bulb quickly. This translates into slower reaction times on this measure for outcome-dependent participants. This is indeed what the study found (see Figure 4.5).

This research provides evidence for the view that motivated involvement with a target can lead to more controlled processing (and hence less stereotypic impressions; Neuberg & Fiske, 1987). More than this, it suggests that the extent to which we are able to go beyond initial, category-based impressions will be dependent upon the interplay between motivational and attentional factors. In sum, once attention is depleted, our ability to systematically process information about others, even if we are motivated to do so, may be diminished.

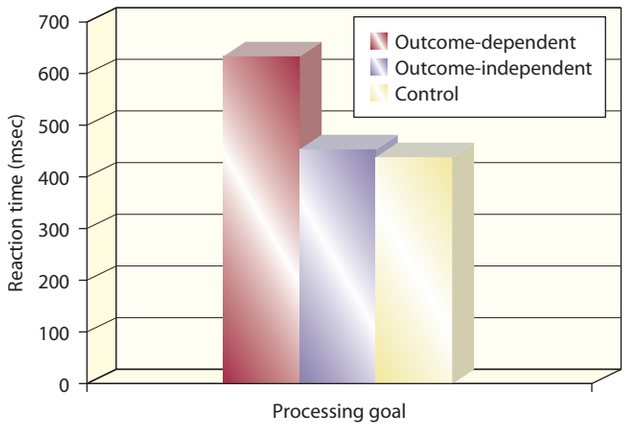

Figure 4.5 *Participants' mean probe reaction times (PRTs) in msecs as a function of processing goal (from Pendry & Macrae, 1994, Study 2).*

Replacing stereotypic thoughts with egalitarian responses As we saw earlier, Devine's (1989) paper provided some initial evidence for the automaticity of stereotype activation. Devine's thesis did not, however, stop there. In her ***dissociation model*** of stereotyping, she argued that automatic and controlled processes may be dissociated. What this means is that automatic activation of a stereotype does not inevitably lead to stereotypic responding.

dissociation model a model that proposes that two different processes can occur independently, and that one does not inevitably follow from the other (e.g., Devine's theoretical model that proposes a dissociation between automatic and controlled processes in stereotyping)

Later research echoes this general sentiment. For example, Monteith (1993) has shown that when people are committed to being non-prejudiced and their behaviour appears to violate these standards, they feel guilty, become self-focused (direct attention towards the self) and direct their efforts at reducing this discrepancy to ensure it does not happen again. Hence, Monteith (1993) found that low-prejudice participants provided unfavourable (i.e., non-stereotypic) evaluations of jokes about gays, but only if they had been made to realize that in an earlier phase of the experiment they had (without realizing it) acted in a prejudiced fashion.

What research of this kind demonstrates is that it is possible to regulate stereotypic responding if (1) we are aware of the possibility of unconscious prejudicial influence, (2) we are sufficiently motivated (here, by virtue of a desire to appear unprejudiced) and (3) we have the required time available to do so (see Macrae & Bodenhausen, 2000). There are several issues of note here. For example, we may not always be aware of the unfelt influence of the stereotype (Bargh, 1999; Wilson & Brekke, 1994). Also, as we have seen, time or processing capacity limitations can impede even the most motivated perceiver (Pendry & Macrae, 1994). It is possible, too, that even if we are motivated to control stereotypic reactions, attempts at control can backfire for an altogether different reason, as we shall now see.

Stereotype suppression: Pushing the unwanted thought out of mind (if not always out of sight) Imagine you have just encountered an elderly woman in the supermarket. She looms large as you enter the fruit and veg aisle, thwarting your speedy passage to secure an aubergine by inconveniently standing right in front of you, consulting her shopping list. 'Dithery old biddy, it must be pension day!' you catch yourself thinking, and then you chastise yourself. You think, 'I really must stop this, she is no more in my way than anyone else, she is just rather older than most'. You try to banish such stereotypic thoughts and proceed to the dairy aisle. There, you encounter another elderly female. She is also in your way, but this time she's carefully weighing up the prices of different cheeses as you wait to extricate the last packet of Parmesan from the depleted shelf above her. How do you react to this second elderly female? Are you successfully able to suppress the elderly stereotype?

stereotype suppression the act of trying to prevent an activated stereotype from impacting upon one's judgements about a person from a stereotyped group

There has been a great deal of interest in precisely this topic: does ***stereotype suppression*** work? The research was stimulated by Wegner's (1994) ironic processes of mental control model. According to Wegner, when we try to suppress unwanted thoughts, two mental processes result. First, the intentional operating process (IOP) begins to search for thoughts that can serve as distractors (to distract us from thinking about the thing we don't want to think about). At the same time, a second, ironic monitoring process (IMP) kicks in, searching for evidence of the unwanted thoughts. In order to identify these unwanted thoughts, the IMP has to hold at some preconscious level the very thoughts one wants to suppress. Here's the important point: the IOP is a cognitively demanding process. To use the terminology we employed earlier, it entails controlled thinking. However, the IMP is thought to operate in an automatic manner.

So what happens next? Because the IMP (the process that spots signs of suppression failure) can operate in the absence of cognitive resources, it is free to run mental riot even when resources are depleted, constantly searching for signs of failure (i.e., of the unwanted thought itself). Recall earlier we learned that constructs that were frequently activated (primed) become more accessible. Well, that is pretty much what is hypothesized to happen here. The unwanted thoughts on which IMP is focusing receive a healthy dose of priming and, without the IOP, become even move accessible. In other words, a ***rebound effect*** is demonstrated. The implication for stereotype suppression is that, under certain circumstances, the more people try to suppress stereotypes, the less successful they will be.

rebound effect where suppression attempts fail; used here to demonstrate how a suppressed stereotype returns to have an even greater impact upon one's judgements about a person from a stereotyped group

Macrae, Bodenhausen, Milne and Jetten (1994) reported a series of experiments that demonstrate this rebound effect. In their first study, participants were asked to write about a day in the life of a skinhead, with a photo as a prompt (the study purportedly investigated people's ability to construct life event details from visual information). Half were told to avoid stereotypic thoughts about skinheads (i.e., suppress stereotype) while writing the passage, half were not. Later, they were shown another skinhead photograph

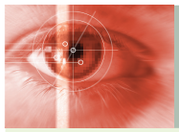

RESEARCH CLOSE-UP 4.3

Automatic and controlled components of stereotypes and prejudice

Devine, P.G. (1989, Study 3). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.

Introduction

In her first study in this series, Devine demonstrated that all participants, whether high or low in prejudice, were equally knowledgeable of the cultural stereotype of blacks. The second study demonstrated that when participants' ability to consciously monitor stereotype activation was prevented, all participants responded in line with the activated stereotype (as we saw earlier). Devine's thesis did not, however, stop there. In her theoretical model, she argues that automatic and controlled processes may be dissociated. What this means is that automatic activation of a stereotype does not inevitably lead to stereotypic responding. When participants have time and motivation to correct for initially stereotypic thoughts, they will do so. Later on in the paper (Study 3), she set out to demonstrate this.

Method

Participants

Sixty-seven white students took part in the study. Participants were divided into high-prejudice ($N = 34$) and low-prejudice ($N = 33$) groups based on a median split of scores on the Modern Racism Scale (MRS; McConahay, Hardee & Batts, 1981).

Design and procedure

The design involved a simple one-way comparison between participants low vs. high on prejudice. Participants were run in small groups. First, they were asked to list as many alternate labels as possible for the social group black Americans (to include slang terms). This served to activate participants' cognitive representations of blacks. Following the label-generation task, they were asked to list their honest thoughts about the racial group blacks, under anonymous conditions. Afterwards, they completed the seven-item MRS.

Results

The proportion of pejorative and non-pejorative labels arising from the label-generation task was computed for each participant. A comparison between high- and low-prejudice participants revealed no significant differences in terms of the proportion of pejorative labels generated in the first phase. Then participants' responses to the thought-listing task were coded in terms of valence (positive or negative) and whether the thought concerned a belief about the group or was instead a trait description. Thus, there were four different kinds of thoughts coded (positive trait, negative trait, positive belief, negative belief).

Analyses of the frequencies of different types of thoughts listed by participants revealed that high-prejudice participants more often listed negative traits than each of the other three types of thoughts (which did not differ from each other). However, low-prejudice participants listed positive belief thoughts more often than the other three types of thoughts (which did not differ from each other).

Discussion

This study demonstrated that low-prejudice participants were able to provide non-stereotypic and egalitarian descriptions about blacks *provided they were given sufficient time to generate these descriptions*. As such, it appears to qualify the rather pessimistic conclusions highlighted by the first two studies in the paper.

It should be noted, though, that there are a few methodological issues associated with this paper. For example, although the first study appeared to show no differences in activation as a function of prejudice level, we cannot be sure activation differences do not exist (perhaps future research using a different method might still detect differences in activation). Allied to this, the MRS measure used in these studies may not be the most sensitive (or indeed the only) way to look at individual differences (as we saw with the Moskowitz et al., 1999, study). Second, the primes used were both category (e.g., *nigger*) and negatively valenced trait (e.g., *lazy*) types, so it is difficult to be sure that it is category priming and not simply semantic priming that is driving the effects. Later research focusing just on category labels (e.g., *blacks*) and neutral semantic associates (e.g., *ethnic*) showed that low-prejudice participants displayed less negative reactions to outgroup-related primes than did high-prejudice participants (e.g., Lepore & Brown, 1997). Finally, the stimulus materials were words, not pictures or real-life interactions. It may be unwise to assume effects will inevitably be similar irrespective of the nature of the prime (e.g., Gilbert & Hixon, 1991). Whilst later research has qualified some of these findings, this remains an extremely influential and widely cited paper in the field.

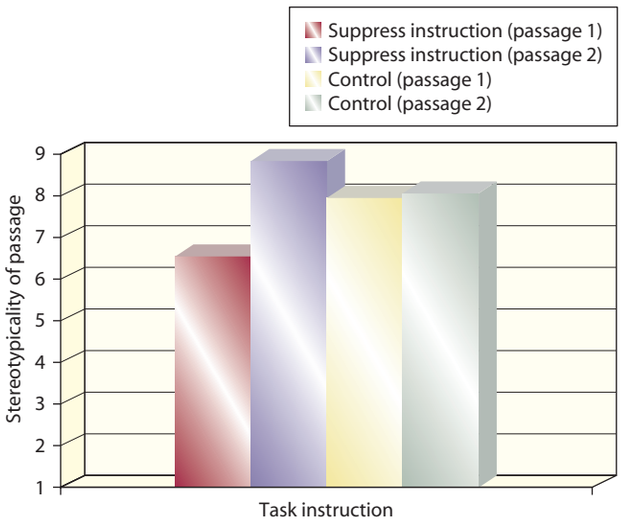

Figure 4.6 *Ratings of passage stereotypicality as a function of task instruction (from Macrae et al., 1994, Study 1).*

and asked to write a second passage. This time no 'skinhead suppression' instructions were given. The researchers hypothesized that if the 'suppression' participants experienced repeated stereotype priming in the first phase, then they might show evidence of a rebound effect in the second phase. As a result, their passages should be more stereotypic in the second phase. This is indeed what Macrae et al. found (see Figure 4.6).

Two additional studies provide further support for this finding. In a second study, the rebound effect was demonstrated in a different way (behavioural reactions). Participants initially suppressing a stereotype in the first phase elected to sit further away from a skinhead's belongings (i.e., where he would presumably return to sit down) in the second phase than participants who were not instructed to suppress. The final study used a lexical decision task to demonstrate that participants who were suppressing a stereotype about a skinhead later showed faster responses to traits related to the skinhead stereotype. This finding suggests that the initial suppression phase resulted in the stereotype becoming hyperaccessible. Later research developed these preliminary findings, in particular by showing that heightened self-focus can cause suppression to occur in a spontaneous fashion (e.g., Macrae, Bodenhausen & Milne, 1998).

Although these results paint a fairly convincing picture, several years and a handful of studies later some caveats are in order that pertain to methodological issues and concerns about external validity (for a review, see Monteith, Sherman & Devine, 1998). Consider the stereotypes used in these early studies. Skinheads (and in other studies, hairdressers, supermodels and construction workers) may not be groups for whom we feel a great need to suppress stereotypes (in comparison to, say, blacks, gays or women). Stereotyping certain groups may not carry the same potential penalties or condemnation. Also, people differ in terms of the extent to which they endorse or avoid stereotyping. Where stereotypes of arguably more sensitive groups such as gays are studied, rebound effects are weakened among participants low in prejudice towards this group (e.g., Monteith, Spicer & Tooman, 1998). However, participants high in prejudice demonstrate the same rebound effects we saw earlier.

Several reasons are advanced for these differences (see Monteith, Sherman & Devine, 1998). It may be that people low in prejudice are more motivated to avoid prejudiced reactions (Fazio & Dunton, 1997); or that they are more practised in trying to rid themselves of stereotypic thoughts. Perhaps they have faster access to replacement (i.e., non-stereotypic) thoughts (e.g., Blair & Banaji, 1996); or they may be more motivated to form individuated impressions of others (e.g., Fiske & Neuberg's, 1990, continuum model mentioned previously). Finally, they may possess a goal state that encourages them to create a particular desirable state of mind (e.g., chronic egalitarian goals, as in the research by Moskowitz et al., 1999, we saw earlier) as opposed to suppressing an unfavourable one (stereotypes). Monteith, Sherman and Devine (1998) suggest that an important goal for future research will be to identify the part such factors play in stereotype regulation. For now, though, the initially gloomy picture about the consequences of suppressing stereotypes seems unwarranted.

The link between social perception and social behaviour is not inevitable Earlier, we saw some intriguing demonstrations of the link between stereotype activation (e.g., priming participants with elderly traits) and behaviour (e.g., participants walking more slowly). This seems quite compelling evidence for the inevitability of stereotype activation. Subsequent research has, however, enabled us to build a more balanced picture. It appears that whilst this effect does often happen, there are several factors which, when present, modify the typical pattern of results.

Many of the studies demonstrating the automatic effects of schema activation upon behaviour fail to take into account the potentially moderating effect of both factors inside the person (perceivers' motivations and goals) and factors outside the person (characteristics of the environment). Macrae and Johnston's (1998) paper neatly fills this gap. In the first study, participants were primed (here, with a trait construct: helping or not helping). As they were preparing to leave the experimental room to move to an adjacent laboratory, the experimenter dropped her belongings upon the floor, including a number of pens. Importantly, in one condition the pens were leaking badly, but in the other condition they were not. What Macrae and Johnston predicted and found was that participants were more likely overall to help following the helpful prime, but only when the pens were normal (helping was high in both help prime and control prime conditions: 93.7 per cent and 68.7 per cent, respectively). But the prime had no effect when the pens were leaky. Apparently, the thought of helping to pick up the pens and getting covered in ink was a strong disincentive to help in both priming conditions (help prime: 6.2 per cent and control prime: 12.5 per cent). In the second study, participants were again primed with the construct of helpfulness. In addition, they were told they were either on time or late. Again, as they got up to leave, the experimenter dropped her belongings, including some pens (none of which was leaking). Whilst participants primed

with helpful were more inclined to help, this tendency was notably decreased for participants led to believe they were running late.

These findings imply that the typical effects of perception upon behaviour are dominated by current processing goals, when the behaviours needed to attain the goals are at odds with those implied perceptually (i.e., even though primed with helpful, the costs of being helpful in terms of getting covered in ink or being late serve to override the effects of the prime). Hence, behavioural control is viewed as a battle between activated schemas and various environmental cues and internal goal states either promoting or inhibiting the occurrence of certain action patterns (Shallice, 1988).

The effects of priming on automatic social behaviour seem also to be eliminated when participants' self-focus is increased (for more on effects of self-focus, see Chapter 5, this volume). Dijksterhuis, Bargh and Miedema (2000) primed participants with the politician stereotype (or did not). Half were seated in front of a mirror (high self-focus), half were not. Later they were all asked to write an essay about nuclear testing. Pre-testing had established that an aspect of the politician stereotype is that they are notoriously long-winded. Hence, the researchers hypothesized that the politician prime would result in longer essays. This was true, but only for participants in the low self-focus condition. Participants seated in front of a mirror did not show the effect.

So why does self-focus diminish the effects of the prime? The researchers argue that self-focus has been shown to activate what are termed *action tendencies* (Carver & Scheier, 1981). The action tendencies that self-focus can make more salient and accessible are certain norms, behavioural standards and goals. Here self-focus effectively serves to prevent the execution of an undesirable behaviour (being long-winded). Under conditions of self-focus, usual effects of perception on behaviour can be eliminated.

So, the situation regarding the seemingly automatic effects of schema activation upon behaviour is rather more elaborate than was first thought. In many cases, though, the effects of stereotypes and other schemas are far from inevitable.

Can schemas change?

So far we have discussed how it is possible for schema-based processing to be overridden by a more considered appraisal of available data. For example, how it is possible for us, once a stereotype has been activated (e.g., blacks are athletic and into gangsta rap; see Johnson, Trawalter & Dovidio, 2000), to avoid its effects in favour of judging a target in a more individuated fashion (this black person is a kind, gentle, classical music-loving person who adores his family). What we have not yet tackled is the extent to which schemas (in the form of stereotypes) can change. So, rather than having to overcome the activated stereotype, is it possible for the schema itself to be modified? It is possible, but that is not to say it is easy.

Recall that schemas, acting as cognitive shortcuts, are a functional way of parsing our social environment. They provide order and predictability. The benefits of having schemas ultimately mean there are often pressures to maintain them (Fiske & Taylor, 1991). Some of these pressures include our desire not to change schemas even in the face of disconfirming evidence; the fact that thinking about a schema strengthens and commits us to holding it even more; and a seeming inability to reconsider if our old schematic beliefs are still applicable in the face of new information (see Fiske & Taylor, 1991). So why bother? As Fiske (2004, p. 156) neatly sums it up: 'the cost of constructing a new schema seems psychologically prohibitive'.

And yet, in certain circumstances, schemas *can* and *do* change. It is, after all, only worth relying upon our schemas when we have reason to believe they are serving us well. Having a schema that is incorrect or inaccurate can lead to errors in judgement and memory. Research suggests that schemas change when they are clearly disconfirmed, when people come across alternative schemas and take the time and trouble to scrutinize unique, individual instances (Fiske, 2004).

Within social cognition, several models of schema (stereotype) change have been proposed (Hewstone, 1994; Weber & Crocker, 1983). In the *bookkeeping* model, schema revision is viewed as a very steady process, whereby each new piece of disconfirming information is logged, leading to a very gradual modification of the original schema. The *conversion* model, however, predicts rather rapid and wholesale change of a schema in response to a large amount of disconfirming information. Finally, the *subtyping* model suggests that subcategories develop when faced with individuals from a category who strongly disconfirm it. This last model is better termed a model of maintenance, as opposed to change, since the formation of such subtypes effectively insulates the pre-existing stereotype from change.

A programme of research by Hewstone and his colleagues has focused closely upon specifying the conditions leading to schema change via the above routes (for a review see Hewstone, 1994; for alternative interventions to reduce prejudice, see Chapter 14, this volume, on intergroup relations). For example, in one study (Johnston & Hewstone, 1992, Study 1) participants (psychology students) received information about a group of physics students. Some of the information was consistent with the stereotype (dress in rather nerdy clothes), some was inconsistent (likes to go out) and some was neutral. Importantly, the pattern of inconsistent information was systematically varied. For some participants, the inconsistent information was concentrated in two out of the eight physics students (so they were really rather atypical); for others, this same information was dispersed across six members (who each only slightly disconfirmed the stereotype). A final intermediate condition distributed inconsistent information across four targets.

What these researchers found was that participants in the concentrated condition rated inconsistent traits as significantly less characteristic of the group than did participants in either the intermediate or dispersed conditions. It is likely, then, that participants in the concentrated condition lumped the two wildly inconsistent individuals into a single atypical subtype and did not incorporate them in their subsequent evaluations of the group. Where there were several examples of individuals who each disconfirmed the stereotype, even if only slightly (as in the dispersed and, to some extent, the intermediate conditions), this inconsistent information was rated as more characteristic of the group.

The take-home message here is that for schemas to change, it is not enough simply to encounter disconfirming information.

Schema change may be more likely when we encounter several people who seem to be just a bit unlike the stereotype, as opposed to encountering just one or two who really seem to disconfirm it greatly. The fit between target members and group is so weak that little or no effort is made to integrate the inconsistent information. It is hence easier to discount these extreme individuals as too atypical to be taken seriously. In sum, stereotypes can change in response to disconfirming information, but an important moderating factor is the manner in which the disconfirming information that should promote change is presented to us.

SUMMARY

In this section, we have seen how it is sometimes possible to exert control over stereotype activation. Moreover, we may still be able to rescue the situation even if stereotypes have been activated, provided we are aware of the potential influence of the activated stereotype, are motivated not to stereotype and are cognitively able to do so. Finally, we have seen how, under certain conditions, it is possible for stereotypes to change. So, the picture may be less bleak than we might have feared.

SUMMARY AND CONCLUSIONS

Now that you have reached the end of this chapter, you are in a better position to provide answers to some of the questions posed at the beginning. Let's revisit them and then recap on what we have covered.

- Why did I assume that the man at the coffee machine in the boardroom was the company director when he was in fact the secretary?
- Why did I assume that Dr Alex James would be male/white?
- Why is it that I expected Albert to be elderly?

Answer: automatic stereotype activation. Categories like gender, race and age are readily activated in the presence of a person from or a name associated with these groups.

- Why did it surprise me to discover that Hilda, my elderly neighbour, had a passion for car maintenance?

Answer: we expect, and often seek out, information that is consistent with our stereotypes. An elderly female's penchant for wielding the spanner violates our well-established expectancies of what little old ladies typically do.

- Why did I take the time to talk to my new black work colleague and find myself subsequently re-evaluating my initially stereotypic impression of her?

Answer: when we are motivated – for example by virtue of needing to get along with someone or because we are low in prejudice towards members of that group – and have the cognitive resources available, we are able to move beyond initial category-based impressions to form more individuated ones.

- Social cognition research has provided us with some important theoretical clues about when and why we engage in automatic versus controlled processing of social information.
- Automatic processes are those that occur without intention, effort or awareness and are not expected to interfere with other concurrent cognitive processes. Controlled processes are intentional, under an individual's volitional control, effortful and entail conscious awareness.
- Often, stereotype activation can occur automatically. Once a category is activated, we can bring into play the knowledge contained within these structures (schemas). Schemas affect how quickly we perceive and interpret available information, and impact on subsequent processes of judgement and memory. They can also impact upon our behaviour, as shown by research into the perception–behaviour link.
- Sometimes we process social information more systematically. We may, under certain circumstances, not activate stereotypes at all. If we do activate them, we may engage in several strategies to avoid responding in a stereotypic way. For example, we may engage in a more complex appraisal of the available information (individuated impression formation), replace stereotypic thoughts with more egalitarian ones or attempt to suppress the stereotype. Under certain conditions, too, stereotypes can change.
- Some researchers, such as Bargh (1999), consider that stereotype activation is more inevitable than we might like. Others, like Devine and Monteith (1999), take a more cautious view, arguing that control appears to be possible, at least some of the time.
- This chapter has used the automatic/controlled distinction as a focus to introduce you to some of the fascinating theoretical questions and research methodologies that typify social cognition research. The research conducted in this area speaks to issues of considerable social importance. Researchers will continue to pose intriguing questions and develop yet more sophisticated ways in which to assess the complex processes that underlie our daily mental life.

Suggestions for further reading

Bargh, J.A. (1999). The cognitive monster: Evidence against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 361–382). New York: Guilford. A spirited and utterly engaging defence of the inevitability of stereotype activation.

Devine, P.G. & Monteith, M.J. (1999). Automaticity and control in stereotyping. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 339–360). New York: Guilford. A slightly more even-handed debate on the same topic.

Dijksterhuis, A. & Bargh, J.A. (2001). The perception–behavior expressway: Automatic effects of social perception on social behaviour. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 33, pp. 1–40). San Diego, CA: Academic Press. This chapter brings together much of the recent literature on this topic and attempts to provide a better understanding of the mechanisms that may underlie the effects.

Fiske, S.T. & Taylor, S.T. (1991). *Social cognition* (2nd edn). New York: McGraw-Hill. This is the classic text on the topic, with extensive coverage of a wide range of issues within the discipline.

Macrae, C.N. & Bodenhausen, G.V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, *51*, 93–120. A thorough, readable overview of the literature.

Monteith, M.J., Sherman, J.W. & Devine, P.G. (1998). Suppression as a stereotype control strategy. *Personality and Social Psychology Review*, *2*, 63–82. A clear, considered and interesting review of the literature.

Moskowitz, G.B. (2005). *Social cognition: Understanding self and others*. New York: Guilford. A welcome new addition, bang-up-to-date, engagingly written and comprehensive in scope.