14

Implicit Impressions

James S. Uleman, Steven L. Blader, and Alexander Todorov

Our knowledge of other people is complex and multilayered. It develops from a thousand strands: a friend's comment ("she's cool"), taste in clothing (all from The Gap), tones of voice (like a fishwife), attitudes (against the death penalty), values (pro-life), facial expressions (crooked smile with overbite), favorite haunts (jazz clubs) and foods (eggplant and garlic), what someone does with us (talk politics, gossip) and others (makes love), roles (sister-in-law), goals (stay sober), style (shy at first), and so on ad infinitum. What we know quickly extends beyond what we can say to how we feel, what we implicitly expect, and how we act toward them and with them, sometimes to everyone's surprise. Much of it is based on forgotten encounters and barely noticed events. It includes the meanings we give these events and observations, sometimes consciously with care and effort, but more often without effort or awareness.

In this chapter, we begin to develop the idea of implicit impressions of others. Implicit impressions are those impressions—or that aspect of all impressions—that do not depend on explicit memory of past encounters or the explicit meanings people have attached to them. They extend beyond verbal descriptions of others to include the intuitive basis of those descriptions that are called into play when we try to put our implicit knowledge into words for the first time. They are what is at work when we notice that someone is acting "out of character," without quite being able to say how we know this or what their character is like. They are what is behind our certainty that we "know" what someone is like without being able to say (or remember) what we know or how we know it. Implicit impressions are preverbal, nonepisodic residues in memory of our observations of, interactions with, and inferences about others. These memorial traces are linked to other people in forms that guide our explicit thoughts, our emotions, and our behavior with them, without having to emerge from the shadows and become explicit themselves.

By definition, implicit impressions are not something one can tell to others. But knowledge about others clearly passes back and forth between the implicit and explicit domain. It moves from implicit to explicit every time we first articulate it—even when it has been clear to others for years in our gestures or feelings or actions. It moves from explicit to implicit when it operates outside current awareness because we either do not remember it or cannot remember it. This explicit-turned-implicit knowledge may be about our own prior verbalizations or about their evidentiary bases that we could have described.

Distinguishing implicit from explicit impressions can be difficult. Sometimes people say one thing about someone but do another, so words and deeds are at odds. Excluding deliberate deception, such cases suggest that both explicit and implicit impressions are at work and are discrepant with each other. But when words and deeds are not in sync, the distinction depends on the memory status of the impression and the information on which it is based. Memories are implicit whenever they affect current behavior (so we know they exist) but have not been, or cannot be, deliberately retrieved. Schachter (1987) described implicit memory as "revealed when previous experiences facilitate performance on a task that does not require conscious or intentional recollections of those experiences," whereas explicit memory "requires conscious recollection of previous experiences" (p. 501). So responses to others that are affected by prior experiences with them, in the absence of explicit reference to or recollection of those prior experiences, are the result of implicit memory. We call such implicit knowledge—which affects how we think, feel, and act toward specific others—implicit impressions.

We want to make several points in this chapter. First, implicit impressions exist, and there is good experimental support for them, including Andersen's work on social-cognitive transference (see chapter 16), and Carlston and Skowronski's work on spontaneous trait inference (STI). Second, implicit impressions affect trait judgments of others. This is clear in Andersen's work and has been demonstrated in research on spontaneous trait transference (STT; Skowronski, Carlston, Mae, & Crawford, 1998), in which STIs affect trait ratings of a communicator rather than the actor. Because STTs are misattributions, they raise the question of how likely it is that STIs will be associated with the correct person, the actor, when other persons are present. New findings from our lab suggest that this is the rule and that STT errors are the exception. Third, implicit and explicit impressions of the same person can be held simultaneously, and their effects can be empirically distinguished. We illustrate how to do this with Jacoby's process dissociation procedure (PDP; see chapter 15), and new findings from our lab. Fourth, implicit impressions are potentially related to a wide range of phenomena that interest social psy-

chologists, including errors in judging how well people know someone, stereotypes, and in-group/out-group perceptions. Speculations and evidence are offered to persuade the reader of this.

Existence Evidence

What evidence is there that implicit impressions exist? The clinical literature contains numerous anecdotes that illustrate implicit memory phenomena in social interaction. Perhaps the earliest and best known is Edouard Claparède's (1911/1951) "curious experiment" with a patient with Korsakoff's syndrome. Although she had been hospitalized for five years, she did not recognize the doctors and nurses she saw every day. "She forgot from one minute to the next what she was told, or the events that took place" (p. 68). "[T]o see whether she would better retain an intense impression involving affectivity, I stuck her hand with a pin hidden between my fingers. The light pain was as quickly forgotten as indifferent perceptions; a few minutes later she no longer remembered it. But when I again reached out for her hand, she pulled it back in a reflex fashion, not knowing why. When I asked for the reason, she said in a flurry, 'Doesn't one have the right to withdraw her hand?' and when I insisted, she said, 'Is there perhaps a pin hidden in your hand?' To the question, 'What makes you suspect me of wanting to stick you?' she would repeat her old statement, 'That was an idea that went through my mind,' or she would explain, 'Sometimes pins are hidden in people's hands.' But never would she recognize the idea of sticking as a 'memory'" (pp. 69–70). One might say that she had formed a negative but implicit impression of Dr. Claparède. But is there more systematic evidence of implicit impressions, especially with nonclinical populations?

Implicit knowledge of various kinds—attitudes, self-esteem, and stereotypes—are familiar phenomena in social psychology by now (e.g., see Fazio, 1986; http://www.harvard.edu/implicit). Greenwald and Banaji (1995) note that their defining feature "is that traces of past experience affect some performance, even though the influential earlier experience is not remembered . . . is unavailable to self-report or introspection" (pp. 4–5). Their literature survey is largely confined to implicit evaluations, including evaluations of the self and of social groups. However, the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), which is at the empirical heart of most of this work, is not limited to assessing implicit associations between pairs of bipolar concepts and evaluative attributes. For example, Greenwald and Farnham (2000, Experiment 2) measured implicit associations with masculine and feminine attributes. Any attributes can be implicitly associated with any concept, or person.

Measuring implicit associations between concepts with the IAT requires that the concepts be explicit, so they can be presented in the test. Implicit associations are detected through their effects on speeded responses to pairs of explicit concepts on the IAT. Implicit impressions may also involve implicit associations. But impressions may be implicit because the concepts that comprise them are implicit, or events' implications are not (or cannot be) recalled. We may not know that we "know" things about other people, yet we still respond to them differently than we would have without some previous forgotten information about them. Even in the absence of explicit concepts, implicit impressions can be revealed by the classic implicit memory phenomena described by Schachter (1987; Schachter & Curran, 2000). Research summarized below demonstrates implicit impressions' impact on trait ratings. But we suspect that they also have effects on emotions, moods, expectations, nontrait descriptions, and behaviors toward others.

One might think that research with the IAT provides evidence of implicit impressions. Although we see no reason why it could not, it has not. All of the research with the IAT to date examines associations with objects, with social categories of people (often ethnic groups), or with the self. None of this work (to our knowledge) looks at implicit associations with specific individuals other than the self.

Lewicki reported a series of studies in which participants either had a brief interaction with one person (Lewicki, 1985) or observed several photos of people (Lewicki, 1986). The person and photos were associated with either positive or negative events, and each had a subtle distinctive physical feature. Participants apparently learned this association between valence and feature nonconsciously, because it affected their responses to a new person who had this feature, even though they reported no awareness of it.

Probably the most extensive body of systematic evidence for implicit impressions comes indirectly from the remarkable research by Andersen and her colleagues (e.g., see chapter 16, this volume; Andersen & Berenson, 2001; Berk & Andersen, 2000; Chen & Andersen, 1999) on social-cognitive transference. This work has shown that rich and complex representations of a significant other can affect a whole array of responses to a stranger who resembles that significant other—responses that range from mood to impressions to expectations and even to interpersonal behavior. These effects occur without any awareness that the representation from the past (which participants previously described) affects current responses to the stranger. Andersen's experimental procedures are elaborately designed to make explicit recollection of the significant other highly unlikely (see especially Glassman & Andersen, 1999). This extensive body of research demonstrates that implicit impressions exist by documenting their effects on a wide variety of reactions to strangers who resemble these significant others.

There is also extensive evidence of implicit impressions created in the laboratory by brief, incidental encounters. In this research, neither social-cognitive transference nor preexisting impressions of social categories or their members are at work. Carlston, Skowronski, and their colleagues developed their savings-in-relearning implicit memory paradigm to investigate whether memory links are established between actors and spontaneous trait inferences (Carlston & Skowronski, 1994; Carlston, Skowronski, & Sparks, 1995). STIs are trait inferences that are made in the absence of intentions to form impressions or infer traits about others. They occur in the process of comprehending trait-implying behaviors or behavior descriptions. For example, most undergraduates reading, "The secretary solved the mystery halfway through the book" spontaneously infer *clever*, even if they are simply reading the sentence as a distracter from a primary task (Winter, Uleman, & Cunniff, 1985) or for a memory test (Winter & Uleman, 1984). Evidence that STIs occur has been provided by a variety of dependent variables, including cued recall, lexical decision reaction times (RTs), and recognition probe RTs (see Uleman, Newman, & Moskowitz, 1996, for a summary). Carlston and Skowronski's work shows that STIs can also function as implicit impressions.

Carlston and Skowronski's procedure was developed to address a controversy about whether STIs are about the actors (e.g., the secretary) or merely about the behaviors. While STIs about behaviors might be interesting, the stronger traditional claim has been that they are about the actor (Winter & Uleman, 1984). If they are not about the actor, they are not implicit impressions. In Carlston and Skowronski's procedure, participants first familiarized themselves with a series of photos of people ("actors") paired with trait-implying paragraphs that described their behavior. Then after intervals as long as one week, they were asked to learn pairs of photos and traits, without any reference to their prior familiarization experience. These photo-trait pairs included the actor photos presented earlier, paired with the traits implied by the paragraphs with which they appeared. These "old photo–implied trait" pairs were easier for participants to learn than either (1) "new photo–implied trait" pairs or (2) old photos paired with other implied traits. That is, there was savings in relearning, in the sense that participants apparently learned pairings of photos and implied traits during the initial familiarization procedure, and then merely had to relearn the "old photo–implied trait" pairs later. Thus learning was easier for old photo-trait pairs. All of this occurred without any suggestion to participants that they should infer traits or form impressions, so trait inferences were spontaneous. Furthermore, the findings established that the STIs were linked to the actors rather than merely the behaviors, because savings occurred for actor-trait pairs, not for behavior-trait pairs.

This research also showed that these savings effects still occur even when learning occurs at least 2 days after the initial exposure, and even when participants show neither recognition memory for the behaviors (Carlston & Skowronski, 1994, Studies 2–4) nor recall of the inferred traits (Carlston et al., 1995, Study 5). That is, the savings occur not only with no explicit reference to the initial familiarization exposure, but also with no explicit memory for the trait-implying behaviors or the traits they implied. This "suggests an expansion of our conception of trait inference or trait knowledge to include implicit as well as explicit knowledge" (Carlston et al., 1995, p. 433). People were unaware both of making the spontaneous trait inferences and of their influence on their subsequent learning.

Evidence That Spontaneous Trait Inferences Are Linked to Actors

Carlston and Skowronski's (1994; Carlston et al., 1995) research shows that implicit impressions can be created in the laboratory and that they are "about" (i.e., linked in memory to) actors rather than merely behaviors. But it also raises the question of why prior research results with a cued-recall paradigm were so equivocal on the latter point (see Uleman et al., 1996). There are many differences between these two lines of research, including instructions (memorize vs. familiarize), the length of the trait-implying descriptions (one sentence vs. paragraph), how the actors were presented (verbally vs. visually), and dependent variables (cued recall vs. savings in relearning). A series of studies using a false recognition paradigm shows that robust actor-STI links can be created with memory instructions, one-sentence descriptions, and visual presentation of actors.

Todorov and Uleman (2002) asked participants to memorize actor photos paired with one-sentence behavior descriptions. All these descriptions implied traits, but some descriptions also made the traits explicit whereas others did not. Participants were then tested with pairs of actor photos and traits, and judged whether the trait had been explicitly presented earlier with the photo, as part of the behavior description. Six studies showed high false recognition of implied (but not presented) traits paired with target actors -higher than false recognition of implied traits paired with other familiar actors, and higher than that of other traits paired with target actors. These effects occurred even when 120 photo-behavior pairs were exposed for only 5 seconds each. And they occurred in the absence of both recognition and recall of the trait-implying behaviors. That is, implicit impressions produced explicit memory errors.

Two of these studies showed that participants went beyond the spontaneously inferred traits and exhibited less false recognition of antonyms of these inferred traits (relative to unrelated control traits) when these were paired with relevant actors. Apparently they were able to use their implicit impressions to more accurately reject traits opposite in meaning from the behaviors'

implications. Furthermore, across all studies and behavior-trait pairs, the strength of explicit trait judgments about the actors from these behaviors was strongly related both to the occurrence of false recognition and to slower response times for correct rejection of implied traits. That is, these false recognition and response time effects were highly correlated with other participants' explicit trait judgments about these actor-behavior pairs. Furthermore, they were unrelated to the probability of generating the traits to the behaviors alone. This is additional evidence that the implicit impressions were about the actors rather than merely the behaviors. Additional studies (Todorov and Uleman, 2003) showed that these false recognition effects occur relatively automatically: under rapid (2 second) exposures, under shallow information processing, and in the face of a concurrent cognitive load at encoding.

This research shows that STIs from single behaviors can be bound to actors in explicit memory, even in the absence of memory for the behaviors on which they were based. These STIs have implications for judgments about other traits' (antonyms') links to actors. Thus evidence of strong, spontaneous actor-trait links is not limited to implicit memory measures such as savings in relearning. This new false recognition measure reveals them under memory instructions, with one-sentence trait-implying behaviors.[1] So there is a reassuring consistency between the robust, spontaneous actor-trait links in implicit memory revealed by the savings-in-relearning paradigm, and the strength of spontaneous actor-trait links in explicit memory revealed by the false recognition paradigm. Both implicit and explicit memory measures provide clear evidence that spontaneously inferred traits are linked to actors, relatively effortlessly and without intentions or awareness.

Thus implicit impressions not only exist, but can be created "from scratch" in the laboratory. STIs are one kind of implicit impression, created from exposure to others performing trait-implying behavior. They lead to savings in relearning (an implicit memory effect) and errors of false recognition (an explicit memory effect). The next question is, what consequences do they have, beyond those already cited to establish their existence?

Spontaneous Trait Transference

Probably the most important, ubiquitous, and apparently obvious consequence of implicit impressions is that they automatically affect intentional impressions of actors who perform the trait-implying behavior. We and several other theorists (e.g., Anderson, Krull, & Weiner, 1996; Gilbert, 1998) have assumed that STIs constitute the first stage of what may later become intentional inferences. However, this sequence has only been assumed. It has not yet been demonstrated empirically. The first empirical evidence for this is

presented in the next section. As background for that evidence, it is useful to take a more detailed look at STT.

It is well established that implicit impressions can affect intentional impressions of people other than those who initially activated the trait construct. Lewicki's (1985, 1986) early work suggested this. Andersen's research on social-cognitive transference (cited above; see also chapter 16) contains many demonstrations that implicit impressions of significant others can affect a variety of responses to strangers, including trait judgments. Further, the clinical literature suggests that such effects are widespread and long-lasting, and Andersen's research shows that they occur among normal participants. This suggests that social-cognitive transference often affects our impressions of strangers.

However, neither Lewicki's nor Andersen's work demonstrates that specific trait concepts, spontaneously inferred from behavior, can affect intentional judgments of someone who is not the actor. Such evidence would be useful support for our claim that STIs and social-cognitive transference are both instances of implicit impressions. It would also provide a method for studying implicit impressions' effects, independent of the content and quality of particular long-term relationships with significant others. The best published evidence comes from Skowronski et al.'s (1998) studies of spontaneous trait transference. STT was first demonstrated by Carlston et al. (1995, Study 4) in a study of savings in relearning. Using the familiarization and savings paradigm described above, trait-implying behavior descriptions were paired not with photos of the actors, but with photos of communicators offering descriptions of other unseen people. In order to make it unmistakably clear that these were not self-descriptions, the descriptions were worded in the third person and referred to someone whose gender did not match that of the communicator in the photo. Savings in relearning showed clear evidence of STT, even though effects were about half the size of STI effects in prior savings studies. That is, familiarization with communicator-description pairs facilitated later learning of photo-trait pairs, even though these traits did not describe those in the photos.

Skowronski et al. (1998, Studies 2–4) then showed that STT affects trait ratings of the communicator. In addition, Study 2 showed that STT is trait specific, in that it does not affect ratings of unrelated traits with the same valence as the implied trait. Study 3 showed that STT occurs even when participants are told that the communicator photos and behavior descriptions have been paired randomly and that the communicators did not know the targets being described. This study was designed to undercut the possibility that STT occurs because participants believe that communicators and targets are similar because they know each other. Study 4 showed STT with more naturalistic stimuli. Participants saw videotapes of communicators telling an off-camera interviewer about either themselves or another person. Then they

rated the communicators on the implied traits as well as traits of the same valence. These ratings provided clear evidence of STT, specific to the implied traits. As in Study 3, the magnitude of STT was smaller than that of STI. Study 4 also showed that STT occurs even when participants are instructed to make intentional trait inferences about the actors in the descriptive paragraphs.

The most remarkable thing about this series of studies is that it shows STIs being erroneously associated with salient others (communicators) and affecting trait ratings of them in spite of all kinds of steps taken to ensure that participants did not believe the trait-implying behaviors characterized them, and even though impression formation was not their goal. STT occurred in spite of mismatched genders and participants knowing that photos and descriptions were randomly paired. Wyer and his students showed that when participants are asked to form impressions of a target person from a conversation about the target, either between two other people (Wyer, Budesheim, & Lambert, 1990) or between another person and the target (Wyer, Budesheim, Lambert, & Swan, 1994), their impressions of the communicators are also affected. But STT occurs without any mention of impression formation and is more reminiscent of findings by Manis, Cornell, and Moore (1974) that attitude-relevant information incidentally affects impressions of the communicator. Skowronski et al. (1998) explained STT as the result of (a) spontaneous trait activation during the familiarization task, followed by (b) erroneous association of the trait with the communicator photo, and then (c) the associated traits' influence on subsequent ratings of the communicator. They contrasted this associative process with the more deliberate processes of intentional impression formation, in which inferences are accurately made about actors from the outset.

This series of studies was not primarily designed to show that implicit impressions affect trait ratings of actors. But one study (Skowronski et al., 1998, Study 3) suggests they might. As noted above, in this study participants learned that communicator photos and behavior descriptions were actually randomly paired. Half the behavior descriptions were worded in the first person, so that participants who forgot about the random pairing might have interpreted these descriptions as self-descriptions and produced STIs. Then on the judgment test, these participants might have unwittingly "read traits into faces" (Hassin & Trope, 2000) without explicitly recalling the randomly paired behavior descriptions. Results showed that the effects of first-person descriptions on trait ratings were comparable to the effects of third-person descriptions, and both effect sizes were similar to other STT effects. So these first-person descriptions may have produced STIs that affected trait ratings. But it is also possible that the trait ratings were based on recalling behavior descriptions. Unfortunately, participants' recall of descriptions and whether they were first or third person were not assessed. Carlston and Skowronski (1994; Carlston et al., 1995) have shown savings in relearning in the absence of both recognition and recall of the behaviors, suggesting this effect may not have depended on behavior recall.

Thus STI effects on trait ratings of actors in the absence of explicit memory for, or reference to, the behaviors are certainly plausible. One study (Skowronski et al., 1998, Study 3) is at least consistent with this idea. In the next major section of this chapter, we describe a new method of examining this issue, and results consistent with our speculations. These studies show that implicit impressions (specifically STIs) do have automatic (uncontrollable) effects on trait ratings.

A Probable Boundary Condition on Spontaneous Trait Transference

Before we consider that work, it is worth taking a closer look at what may be a boundary condition on STT. In all the STT studies, the only photos presented were photos of communicators. Does STT depend on the presence of only the communicator's photo? Would STIs become associated with communicators if photos of both actors and communicators were present during familiarization with or memorization of the trait-implying materials? Although we have no direct evidence on this question, new findings from our lab suggest that STT would disappear or be attenuated.

Todorov and Uleman (in press) used the false recognition paradigm described above to see whether STIs would become associated with any facial photo or are uniquely associated with the actor's face. In three studies, single trait-implying sentences were presented simultaneously with the faces of two people of the same age and gender, and participants were asked to study them for a subsequent memory test. The actors were named, and one photo was identified with the actor's name and the other (control face) with a different name. In the subsequent false recognition test, the actor's face produced significant false recognition of implied traits compared to the control face. This suggests that STIs are more strongly associated with actors' faces, if they are presented, than with other faces or persons present. However, this result could have occurred because the second face received less attention. So in Study 4, two trait-implying sentences and two facial photos were presented simultaneously, with each sentence's actor identified with a different face. This ensured that both faces (as well as both sentences) would be attended to equally. In the subsequent false recognition test, actors' faces produced more false recognition of implied traits than control faces, even though the same faces represented actors and controls for different traits.

This suggests that STT might disappear or be attenuated if actors' faces were presented along with the communicators' faces, using the familiariza-

tion instruction and savings-in-relearning paradigm of Carlston and Skowronski (1994; Carlston et al., 1995). STT may only occur when there is no visual representation (e.g., a photo) of the actor. However, the necessary research—with savings in relearning and trait ratings as dependent variables—remains to be done.

Overall, research on STT strongly suggests that STIs can have implicit effects, that is, effects that make no reference to and do not depend on recalling the behavioral evidence on which they are based. The work by Todorov and Uleman (2002, 2003, in press) provides good evidence that STIs based on single behaviors are linked to appropriate actors' faces, even when they are present at the same time as others' faces. So it is clear that single behaviors spontaneously produce implicit impressions of actors, in the form of STIs. These impressions, made without awareness and without explicit memory of their behavioral bases, may then have automatic effects on subsequent explicit judgments of the actors. In the next section, we describe evidence that directly supports this idea.

Implicit Impressions' Automatic Effects on Trait Judgments

Cognitive processes can be automatic in a variety of ways, as Bargh (1994) pointed out so clearly. Perhaps the most important sense of "automatic" is being uncontrollable (rather than taking place without awareness, being unintentional, or being extremely efficient). Cognitive control is central to a host of social judgment and self-regulatory phenomena (e.g., Wegner & Pennebaker, 1993). Having an implicit impression of someone should affect explicit judgments, as the STT research (above) suggests. The central question we address here is, how controllable are these effects? Of course, if one is not aware of the implicit impression, its influence is necessarily (and perhaps uninterestingly) uncontrollable. But what if one is made aware of implicit impressions and their potential effects? Can these effects be controlled? If so, how much control is possible?

The last two decades have seen an explosion of research in social psychology on automaticity, where the dominant question has been, "Is it automatic or controlled?" Wegner and Bargh (1998) even argued that most of the classic studies in social psychology, beginning with Festinger and Carlsmith's (1959) dissonance studies, are classics in part because they demonstrate how little control people have over their own social behavior. The list of social phenomena—cognitions, evaluations, and behaviors—that are wholly or partly a function of automatic processes continues to grow, to the point that the very existence of freely acting agents seems to be in doubt (e.g., Bargh & Ferguson, 2000).

Yet even as the domain of automatic social behavior has grown and that of controlled processes has shrunk, it has become clear that the basic question of whether particular social behaviors are automatic or controlled poses a misleading dichotomy. Most phenomena of interest to social psychologists are governed by a combination of automatic and controlled processes. One more demonstration that phenomena we used to believe were controlled (or intentional or volitional) are also affected by automatic processes will not change that. Nor will it demonstrate that these phenomena are entirely automatic, or that volition plays no role. Furthermore, as long as we think of automaticity and control as dichotomous and mutually exclusive, and seek more and more evidence of automaticity, control is in danger of being defined by default as merely whatever is not automatic, however automatic is defined. In such a framework, the amount of control depends on the kind of automaticity being considered. Therefore it is important to adopt a more affirmative definition of control and a theoretical framework that includes both automatic and controlled processes, operating at the same time.

Wegner and Bargh (1998, pp. 464–465) offer a typology of the ways automatic and controlled processes can combine, including operating in parallel, one launching the other, one overriding the other, and one transforming into the other. When governance of a phenomenon is not passing back and forth between automatic and controlled processes sequentially, it is shared by both of them. Automatic and controlled processes operate in parallel, or in tandem (when brief enough sequences can be isolated). In both of these cases, the appropriate question is not whether something is automatic or controlled, but how much is automatic and how much controlled? Therefore, we need ways of estimating the relative contributions of automatic and controlled processes to social phenomena. Then we can begin to ask about the conditions that affect the size of the simultaneous automatic and controlled contributions, and move beyond the simple but relatively uninformative question of whether something is automatic at all.

The Process Dissociation Procedure

Fortunately, when mental phenomena depend on the effects of observable prior events, there is a model and a method that can disentangle the automatic from the controllable processes that are based on those events: Jacoby's PDP (see Jacoby, 1991; Jacoby, Toth, & Yonelinas, 1993; Jacoby, Yonelinas, & Jennings, 1997; chapter 15, this volume). The PDP seems to us ideally suited to get at controlled processes because, in a remarkably clear and straightforward operationalization of the idea of personal control, it simply asks participants to control the influence of prior events on their current responses.

To be more precise, the PDP asks participants to include the influence of prior events on their responses on some trials (the inclusion trials) and to exclude those influences on other trials (the exclusion trials). The difference between performance on inclusion and exclusion trials provides the estimate of control. Then (if it can be assumed or demonstrated that controlled and automatic processes operate independently of each other) a few simple equations allow calculation of the influence of automatic processes. The only other requirement for using this method is that there be some criterion of performance accuracy, that is, a way to distinguish between right and wrong answers. The basic ideas are simple and straightforward, but operationalizing them is more complex. Perhaps the best way to understand them is to work through a concrete example.

Word stem completions (e.g., tri_ _ _) require finding words in semantic memory that fit the cues presented. It is well known that a recent exposure to a word that can complete the stem makes that completion more likely, even in the absence of any explicit memory for that exposure. Two processes are at work here. A word stem may be completed by explicitly recalling recently seen words or by just waiting for the word (e.g., *trials*) to pop into one's mind when one looks at the stem. How much can people control the processes involved in word stem completion? More generally, can we assume that an explicit memory task such as recall is completely controlled because it is deliberate? Or might explicit recall be influenced by automatic processes, such as those that bring a word to mind without effort? Likewise, can we assume that asking people to simply let words come freely to mind will tap only automatic processes and prevent them from occasionally trying deliberate recall? Jacoby (1991) argues that there are few process-pure tasks, tasks that tap purely controlled or purely automatic processes. Most tasks depend on a combination of controlled and automatic processes. The way to separate these processes within any task (such as word stem completion) that depends on prior exposure is to (a) define controlled processes as those that people can control and automatic processes as those they cannot; (b) estimate control by examining the performance difference between conditions when people perform the task by intentionally including that prior information versus intentionally excluding it; and (c) use a set of simple assumptions to estimate the importance of automatic processes, once controlled processing has been estimated.

Jacoby et al. (1993, Experiment 1b) examined the effect of divided attention at exposure on subsequent word stem completions. In the first phase of the study, half the participants (in the full attention condition) read 32 words out loud, presented on a computer screen, and tried to memorize them for a subsequent memory test. The other half of the participants (in the divided attention condition) did two tasks at once, reading and listening. They listened to a tape-recorded list of digits and pressed a key whenever they heard three odd digits in a row (e.g., 1, 7, 3). At the same time, they read words out loud on the computer screen, but tried to not let this reading interfere with their digit-monitoring task. In the second phase of the study, all participants completed word stems created from the words they had read earlier. If the word stem "appeared in green, they were to use it as a cue to help them remember a word that was presented earlier. . . . If they could not think of an old word, they were to complete the stem with the first word that came to mind." If the word stem appeared in red, they were to use it "as a cue for remembering words presented earlier but . . . complete those stems with a word that was not presented earlier" (p. 145). So the green trials were inclusion trials, and red trials were exclusion trials.

Without any concurrent cognitive load at encoding, the probabilities of completing the word stems with an old word in the inclusion condition and the exclusion condition were .61 and .36, respectively. With the concurrent digit detection task, these probabilities in both conditions were .46 (see Jacoby et al., 1993, table 1). Under no load, subtracting the exclusion probability from the inclusion probability gives .25 (= .61 − .36) as the estimate of controlled influences (*C*) of prior exposure to the words. That is, controlled processes increased the probability of word stem completions by .25. The maximum possible control for any task is 1.00, so this represents 25% of the theoretical maximum. Thus under no load, participants had some control but not a lot, relative to the maximum theoretical possibility.

Under load, this estimate of control is 0.00 (= .46 − .46). That is, the concurrent load at encoding completely eliminated controlled processes' contribution to word stem completions. (The fact that load reduced controlled processes to 0.00 is not particularly significant. Less load would simply have reduced control less.)

What about the influence of automatic processes (*A*)? Based on the assumption that *C* and *A* are independent, one can calculate *A* as the exclusion probability divided by (1 − *C*). Under no load, this was .47 (= .36/.75); under load, it was .46 (= .46/1.00; see Jacoby et al., 1993, table 2). That is, the influence of automatic processes was not affected by concurrent load. This result is as we would expect, if *A* truly reflects automatic processes, because automatic processes should be unaffected by cognitive load.

Using the Process Dissociation Procedure to Estimate Implicit Impressions' Automatic Effects

How can the PDP be used to estimate implicit impressions' automatic effects? Several methods are conceivable, but we (Uleman & Blader, 2001) adapted the procedures of Skowronski et al. (1998) to illustrate one possibility. Recall that Skowronski et al. asked participants to rate the personality traits of peo-

ple shown in photos that they had previously seen paired with trait-implying behaviors. So memory of prior exposures (both explicit and implicit) is clearly involved. But unlike effects of exposure to words on word stem completions (as in Jacoby et al., 1993), these are effects of exposure to behavior-photo pairs and the spontaneous inferences that are drawn from them, all of which can affect subsequent trait ratings of these photos. These effects might be completely mediated by implicit impressions (STIs), as suggested by Carlston and Skowronski (1994; Carlston et al., 1995) and Skowronski et al. (1998). They showed that STI occurs, and its effects do not depend on explicit memory for the behaviors. On the other hand, under other conditions, these effects might also be mediated by explicit recall of the behaviors and any trait implications intentionally drawn from them, in response to the trait-rating scale. The PDP model describes how to estimate the effect of both automatic and controlled processes on trait ratings. The automatic processes are the consequence of implicit impressions because, as prior research shows (Uleman et al., 1996), participants are unaware of making STIs. Being unaware of them, they cannot control their effects.

In order to apply the PDP to this phenomenon, two problems had to be solved. First, we had to work out how to obtain trait ratings of the photos under inclusion and exclusion conditions. Second, because there are no objectively correct or incorrect trait ratings of photos of strangers, a way of defining the proportion of correct responses had to be developed. We (Uleman & Blader, 2001) tried the following solutions, in two studies. Participants studied behavior-photo pairs for a subsequent memory test (rather than under Skowronski et al.'s, 1998, familiarization instructions, to maximize the parallels with Jacoby et al., 1993). Just as in the Carlston and Skowronski studies described earlier, participants viewed photos paired with behaviors. Then they got another booklet containing only photos, some of which had been paired with behaviors in the first part of the study and some of which were new. They were asked to examine the photos in this second booklet to form impressions of the pictured people and to rate each one on three traits.

However, the instructions for the rating task were varied to create the inclusion and exclusion conditions. In the inclusion condition, participants were told,

> In the earlier memory task, the pictures were paired with descriptions of behaviors. In every case, those behavior descriptions were about the person in the photo. They were provided by that person. Since the behaviors are about the person in the photo, they're informative and will make it easier to form an accurate impression of him or her. So try to remember the behaviors and then use them as you form your impressions from the photo.

In the exclusion condition, they were told,

> In the earlier memory task, the pictures were paired with descriptions of behaviors. In every case, those behavior descriptions were NOT about the person in the photo. They describe someone else, and do not describe the person in the photo. (We checked with the people to be sure.) We paired them with photos randomly. Since the behaviors are NOT about the person in the photo, they're potentially misleading and may make it harder to form an accurate impression of him or her. So try to remember the behaviors and then ignore them, so you're not influenced by them as you form your impressions from the photo.

In other words, in the inclusion condition participants were asked to remember and use the information they received earlier to form an impression of the person (if the person was among those paired with a behavior in the first part of the study). In the exclusion condition, participants were asked to remember the prior information but then intentionally avoid using it in impression formation, just as participants in the Jacoby et al. (1993) study were asked to remember the words presented earlier and avoid using them to complete the stems.

These instructions were designed to produce an intention to use the earlier trait inference in the inclusion condition and an intention to not use the trait inference in the exclusion condition (rather than simply producing no intention in the exclusion condition). Thus, trait judgments in the inclusion condition should be a result of both controlled (recollection) and automatic processes. However, in the exclusion condition, any influence of prior exposure to the photo-behavior pairs should be the result of uncontrollable (i.e., automatic) processes, because participants exclude the influence of prior exposure as completely as they can.

The proportion of correct responses for "old" (previously seen) photos was defined for trait ratings relative to how the photos were rated in the absence of any prior exposure to them. That is, ratings of those photos on the traits by participants who had never seen the photo before served as the baseline. In the trait-rating phase of each study, 12 of the photos were new to half of the participants, and 12 were old. Participants rated each photo on the same three trait types: the one implied by the behavior, a trait of congruent valence, and a trait of incongruent valence (as in Carlston & Skowronski, 1994). Ratings were averaged among participants for whom the photos were new, creating unique base-rate estimates for each photo of how it would be rated without prior exposure to behavioral information. Among participants for whom the photo was old, a score of 1 was assigned if their rating was above this base rate, and a score of 0 otherwise. Averaged across old photos, this score gives the proportion of photos rated above the base rates, that is,

the proportion correct in the sense that prior exposure to behavior-photo pairs increased trait ratings—just as prior exposure to words increased word stem completions.

In both studies (Uleman & Blader, 2001), we varied the delay between exposure to behavior-photo pairs and trait ratings of the photos. A third of participants had no delay; a third had a 20-minute delay; and a third had a 2-day delay. We predicted that delay would significantly decrease the influence of controlled processes on the trait ratings, but not affect the influence of automatic processes as much (just as cognitive load did not affect automatic processes in the Jacoby et al., 1993, study). The larger objective of these studies was to show that this adaptation of the PDP and Skowronski et al.'s (1998) procedures produces results that are theoretically reasonable and useful for studying the duality of automatic and controlled processes in implicit impressions.

In the first study, we were concerned that it would be difficult for participants to switch back and forth between inclusion and exclusion instructions, as they rated the 12 old photos. So they were given either inclusion or exclusion instructions for all the ratings. This made it impossible to calculate *C* or *A* for each participant, because we did not have both inclusion and exclusion performance for each one. But we were able to calculate *C* and *A* for each critical trait, performing analyses with stimuli as the unit of analysis.

Figure 14.1 shows the values of *C* and *A* as a function of delay. (Values of *A* were adjusted by subtracting the chance probability of .5 that would result if prior exposure had had no influence; see Jacoby et al., 1993, for a discussion of how to treat neutral base rates.) Estimates of *C* were significantly above chance when there was no delay ($p < .001$) and a short, 20-minute delay ($p < .05$). But there was no evidence of significant controlled processing at the long, 2-day delay. On the other hand, estimates of *A* were significantly above chance levels at no delay ($p < .001$), at the short delay ($p < .01$), and at the long delay ($p < .05$).[2] Both *C* and *A* showed significant linear trends over delay ($ps < .01$).

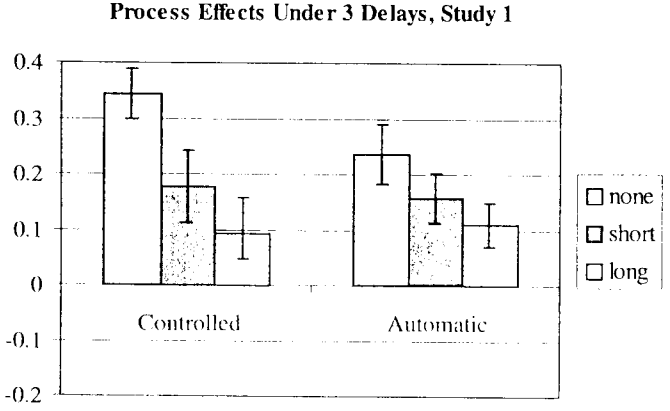

Figure 14.1 Estimates of the magnitude of controlled (left) and automatic (right) effects of prior exposure to behavior-photo pairs on subsequent trait ratings of those photos, under no delay, 20-minute delay, and 2-day delay. Inclusion-exclusion instructions between-Ss; traits as unit.

These estimates of *C* and *A* fit our expectations. First, at brief delays, prior exposure to pairs of trait-implying behaviors and photos do influence trait ratings through both controlled and automatic processes. (It may seem surprising that automatic processes play any role at all after no delay. But with many pairs and the resulting delay between initially seeing a photo and rating it on traits later, some delay and imperfect explicit memory are inevitable.) Second, after 2 days' delay, only automatic processes continued to affect trait ratings. This is consistent with results noted above (Carlston & Skowronski, 1994; Carlston et al., 1995) showing that after 2 days' delay, participants showed neither recognition nor recall of the trait-implying behaviors, yet implicit impressions still had effects.

In Study 2, we directly addressed the question of whether participants could switch between inclusion and exclusion instructions. This study was similar to Study 1 except that participants rated a block of six old photos under inclusion instructions and another block of six old photos under exclusion instructions. Blocks and instructions were counterbalanced, so that some participants had inclusion instructions first and others had exclusion instructions first. Then they were told that the other instruction applied to the next block of trials. Block and order had no effects. Figure 14.2 shows the values of *C* and *A* as a function of delay, again using stimuli as the unit of analysis (for comparability). Estimates of *C* were significantly above chance for the two shortest delays ($ps < .05$), but not at the 2-day delay. On the other hand, automatic processes had effects at the two shortest delays ($ps < .001$) and also at the 2-day delay ($p < .01$). Although *C* showed a main effect and linear trend for delay ($ps < .03$), *A* showed neither ($ps > .15$). Analyses with participants as the unit of analysis showed the same pattern. Thus, essentially the same effects occurred when inclusion-exclusion instructions were varied between subjects and within subjects (although delay no longer affected *A*). This suggests that participants were able to switch decision processes between blocks with different instructions.

More important, both studies (Uleman & Blader, 2001) showed that a 2-day delay had the predicted effect of reducing controlled processing to insignificance, while the automatic effects of implicit impressions on trait ratings persisted. Thus implicit impressions have automatic effects on impressions that people deliberately form of others. And at least with this paradigm, these automatic effects of implicit impressions persist longer than the controlled

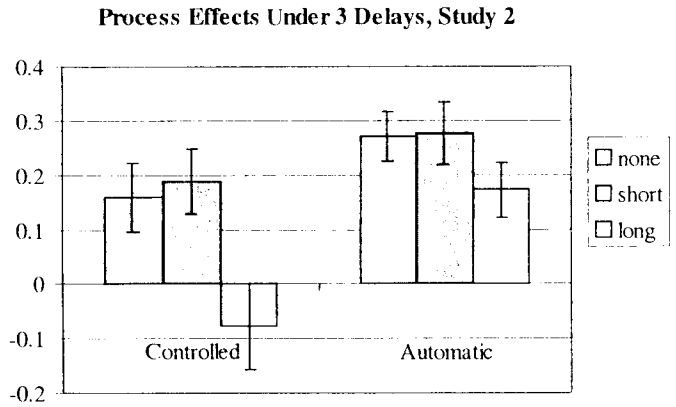

Figure 14.2 Estimates of the magnitude of controlled (left) and automatic (right) effects of prior exposure to behavior-photo pairs on subsequent trait ratings of those photos, under no delay, 20-minute delay, and 2-day delay. Inclusion-exclusion instructions within-Ss; traits as unit.

effects do. At short delays, both implicit impressions and controlled processes had significant effects.

Without something like the PDP and its ability to separately estimate automatic and controlled processing in tasks that engage both, we could not have demonstrated any of this. These results provide another example of why tasks or behaviors should not be called automatic simply because they are influenced by automatic processes—controlled processes may also be operating at the same time. Participants could control effects of prior exposure to the photo-behavior pairs, at least to some degree for shorter delays. Likewise, demonstrating control does not rule out the influence of automatic processes. Prior exposure had automatic effects—the effects of implicit impressions—in all conditions. It takes something like the PDP to let us estimate the contribution of each process type when they both contribute to performance at the same time.

Our adaptation of Jacoby's PDP to study implicit impressions separates their effects from controlled effects. It also takes us beyond the simple question, "Are effects of prior exposures to photo-behavior pairs automatic or controlled?" and allows us to ask, "How automatic and how controlled are they, and under what conditions?" It seems to us that at this point in the development of our knowledge of social cognitive processes, these are the really interesting questions. Even without the present demonstration, we would have guessed that, at short delays, the automatic effects of implicit impressions are accompanied by controlled effects. The challenge was to find a way to demonstrate this and generate estimates of both. Now that we have done this, we can move on to more interesting questions. How do the various encoding, storage, and retrieval conditions that arise in social life affect the automatic and controllable impact of incidental exposure to actors behaving diagnostically? How do individual and cultural differences—which clearly influence the encoding, storage, and retrieval of information about others—affect these automatic and controllable processes? How much can the automatic effects of implicit impressions be brought under control, by telling participants about them and their likely effects? Or would such instructions do nothing more than engage participants' theories about these effects, leading not to control but to correction of automatic effects (as in Wegener and Petty's, 1997, flexible correction model)? This is the kind of research agenda that will move us into the next generation of research on automaticity and control in social cognition.

Other Features of Implicit Impressions

Years of research on STIs have taught us a great deal about implicit impressions, because STIs are a kind of implicit impression. So a brief review of more recent research on STIs is in order here, because it adds to our understanding of implicit impressions.

Implicit impressions should arise not only from printed verbal descriptions of behavior, but also from observing behavior. Skowronski et al. (1998, Study 4) demonstrated STT with videotaped rather than printed stimuli. But the videotaped communicators verbally described behaviors, so the trait-implying behaviors were not enacted. More recently, Fiedler and Schenck (2001) reported evidence for STI from single-frame silhouettes of action. This is an important finding because it provides the first direct evidence that observed behaviors can prompt STI and, by implication, implicit impressions.

Implicit impressions include more than trait inferences. For example, Winter et al. (1985) and Uleman and Moskowitz (1994) showed that "gist cues" (that summarize the action without implying traits) are just as effective as trait cues in the cued recall STI paradigm. That is, people spontaneously encode other aspects of behavior, in addition to their trait implications. Uleman and Moskowitz found the effectiveness of trait cues and gist cues was positively correlated across participants, suggesting that drawing these implications are not mutually exclusive, and that those who infer traits also infer gists. Fiedler and Schenck (2001, Study 2) used the linguistic category model to further explore these processes. They asked participants to verify "whether the indicated behavior could be seen in the preceding picture" (p. 1541). These behavior descriptions were either direct action verbs (DAVs, e.g., *to feed, to hit*), interpretive action verbs (IAVs, e.g., *to nourish, to chastise*), or nontrait adjectives (ADJs, e.g., *sacrificial, severe*). IAVs (which specify goals)

were verified most quickly, suggesting that goals are spontaneously inferred most readily. In addition, subsequent trait identifications (e.g., *caring, merciless*, in a task that gradually revealed perceptually degraded trait terms) were quickest when they were preceded by DAVs, suggesting that DAV verification mediates STIs more than the other linguistic categories. All of this indicates that STIs are not the only spontaneous inferences prompted by descriptions or observations of behavior. Future research should examine the breadth and abstractness of these inferences and how they contribute to the implicit impressions that people form.

Implicit impressions are affected by priming, and can act as primes themselves. When behaviors have ambiguous trait implications, prior activations of trait concepts affect STIs. Newman and Uleman (1990) had participants read such sentences for a subsequent memory test. For example, "Molly would not take 'No' for an answer" can imply that she is either persistent (+) or stubborn (−). Priming with synonyms of these traits (as well as memory for the primes) affected which trait cues were more effective for cued recall. Ferguson and Bargh (2001) found that affective priming, in which the primes have no semantic relationship to the traits or behaviors, produces assimilative STIs. They used the same behaviors and trait cues as Newman and Uleman, but their primes consisted of positively and negatively valenced nouns, such as *sunshine* and *disease*. Finally, implicit impressions can be primes. Moskowitz and Roman (1992) showed this, as did Stapel, Koomen, and van der Pligt (1996). Stapel et al. demonstrated that STIs, as primes, can produce either assimilation or contrast effects, depending on whether the STIs refer to the abstract trait concepts (assimilation) or to particular actors (contrast).

Errors and Biases in Implicit Impressions

At least three decades of research have documented the errors and biases that affect intentional impressions (e.g., Gilbert, 1998), and there is no reason to believe that implicit impressions are immune from these errors and biases. But there may be some errors that are unique to implicit impressions.

Spontaneous Trait Transference

Carlston et al. (1995, Study 4) and Skowronski et al. (1998) have identified one error that implicit impressions can produce: STT. As previously described, STT occurs when particularly vivid communicators (pictured in photos or videotape) describe an absent (not pictured) person in trait-implying terms. Under such circumstances, the descriptions' trait implications are transferred to or associated with the communicator. Carlston et al. (1995) found that

STT did not occur in the savings paradigm when participants deliberately formed impressions of the communicator or target; it only occurred under familiarization (spontaneous) instructions.

However, Skowronski et al. (1998) uncovered a more complex picture. They found (in Study 1) comparable trait transference under familiarization instructions and instructions to avoid STT by forming intentional impressions, when savings were measured after a brief delay. However, after a 2-day delay, trait transfer was evident only under familiarization instructions, but not among participants forming intentional impressions. In Study 4, with trait ratings rather than savings the dependent variable, they found that STT was unaffected by familiarization versus intentional impression formation instructions.

These findings suggest that trait transference—misattributing implied traits to the wrong person—is more likely under spontaneous than intentional inference conditions. They also suggest that STTs have longer-lasting effects than intentional inferences under otherwise comparable conditions. Finally, it is important to remember the boundary condition for STT suggested above. STT may only occur when there is no visual representation (e.g., a photo) of the actor.

Repetition and Apparent Truth

If we hear something false presented often enough, are we more likely to believe it? There is good evidence that we are (Begg, Anas, & Farinacci, 1992; Hasher, Goldstein, & Toppino, 1977), whether these things are presented as facts or opinions (Arkes, Hackett, & Boehm, 1989) or are presented in a visual or an auditory mode (Bacon, 1979), even when we are told that some items are repeated (Bacon, 1979). This frequency-validity effect can persist for weeks (Bacon, 1979; Hasher et al., 1977). Similarly, making the same decision repeatedly increases confidence in that decision (Einhorn & Hogarth, 1978). Expressing the same attitude repeatedly increases its accessibility—one of the components of attitude strength—and hence increases our confidence in it and the likelihood that we will act on it (Fazio, 1995).

One basis for these effects is that explicit thoughts or events that are repeated are also more familiar, and a familiarity-validity heuristic tells us (arguably with some basis in reality) that what is familiar is more likely to be true. However, Whittlesea (1993) has shown that more frequent prior exposure is not the only basis for greater feelings of familiarity, and that other sources of perceptual or conceptual fluency can affect perceived familiarity, often without participants' awareness. Thus, although prior frequency of exposure can be a reliable guide to truth value, using perceived familiarity to judge truth value can also lead to systematic errors. In the domain of impres-

sion formation, Gill, Swann, and Silvera (1998) examined the "representational richness" of people's impressions of others, as one basis for people's widely documented overconfidence in impressions. They showed that representational richness (the extensiveness and integration of person information) is one determinant of judgment fluency, and that judgment fluency increases confidence. Furthermore, other events unrelated to richness (e.g., priming in Study 3) can affect confidence, presumably by affecting judgment fluency.

All of this research (and more) that shows repeated exposure increasing confidence and judged validity uses repetitions of explicit information. But it raises the interesting possibility that the same effects might occur with implicit information. Imagine someone forming an implicit impression of a person (face) on one exposure to a face-behavior pair, and then seeing that face repeatedly without any trait-implying behavior. If the face has become associated with the implicit impression, mere repetition of the face without any explicit repetition of the behavior or its implications should produce greater confidence in trait ratings of the person, in much the same way that repeated exposure to an attitude object increases the accessibility of the attitude toward it. Thus, simple repeated exposure to others about whom we really know very little should lead to increasing confidence that we know them well. Research is needed to test this basic prediction. If it were supported, one could then go on to investigate what conditions promote or impede people's control of this effect, once they are informed of it. Such research would be most informative if it used the PDP adaptation outlined above.

Implicit Impressions, Stereotypes, and In-Groups and Out-Groups

Implicit impressions are about individuals; implicit stereotypes are about groups. Now that we know how to create implicit impressions, it should be possible to create implicit stereotypes—about novel groups with novel content—rather than relying exclusively on existing stereotypes to discover how stereotypes are acquired, change, and function. Crawford, Sherman, and Hamilton (2002) have laid a solid foundation for such research with three impressive studies of implicit stereotype formation, using STT. Participants familiarized themselves with behaviors paired with photos of people from two groups. Results showed more STT (savings in relearning) between members of high-entitativity (e.g., high similarity or coherence) groups than between members of low-entitativity groups or unrelated individuals. In addition, reading times in Study 3 provided evidence of online formation of stereotypes of the high- but not the low-entitativity group. Thus these results demonstrate the development and operation of implicit stereotypes of groups.

At least four other interesting stereotypes and in-group/out-group phenomena are associated with implicit impressions. First, descriptions of actors with ethnic names, who perform behaviors that imply stereotype-consistent traits, spontaneously activate those traits more readily than either the names or behaviors. Howard (2000, Experiments 1 and 2) used a lexical decision task in a series of studies of STIs among Euro-American undergraduates. Participants read sentences and responded to a lexical decision probe immediately afterward. Some of the behaviors implied traits consistent with stereotypes of African Americans (e.g., athletic) or Asian Americans (e.g., shy), and some implied traits consistent with stereotypes of Euro-Americans (e.g., analytical). Some of the actors had typical African American names (e.g., Jamal), while others had typical Asian American or Euro-American names. Lexical decisions for relevant traits were fastest when the actor's ethnicity was consistent with the stereotype-implying behavior. There was no evidence that mismatched names or behaviors activated stereotype traits.

These studies suggest two conditions in which stereotype activation may be less inevitable than many think (e.g., Bargh, 1999). If the eliciting stimulus is weak (e.g., ethnic names rather than faces), spontaneous stereotype activation might not occur. And if it is accompanied by inconsistent information, spontaneous stereotype activation might be blocked. Both of these possibilities need to be examined by further research.

A second phenomenon was uncovered by Todorov, Gonzalez, Uleman, and Thaden (2004). Mismatches between actors and behaviors in terms of gender stereotypes seemed to inhibit implicit impression formation. Using the false recognition paradigm described above, they looked at false recognition of implied traits (in photo-trait pairs) as a function of the trait's gender stereotype and the actor's gender. Relative to pairs with randomly chosen actors, there was higher false recognition of implied traits when they were paired with the actual actor than otherwise, replicating the basic finding described above. But within this result, the false recognition effect was smaller for gender-stereotyped (vs. neutral) traits when the actor was of the other gender. It is not clear from these results whether the mismatch between actors' and traits' genders led to weaker trait activation, weakened actor-trait binding, affected recognition judgments directly through the plausibility of the pairs, or all three. Perhaps the most interesting possibility is that mismatched actors and behaviors prompted spontaneous situation inferences rather than trait inferences. But more research is needed to clarify this phenomenon, the conditions under which it occurs, and the processes that underlie it.

Banaji, Hardin, and Rothman (1993) demonstrated implicit gender stereotyping with a very different paradigm. They showed that the activation of a gender-stereotyped trait (dependence or aggression) primed participants' intentional impressions, formed from ambiguous behaviors of Donald or

Donna, but only when the genders of trait and actor matched. They called this effect

> a previously undocumented effect of implicit stereotyping which . . . we have labeled social category applicability. Social category markers such as race, gender, age, social class, and disability may function like magnetic fields to attract and repel previously encountered stereotyped information[']s impact] on judgment (a) when such information is extraneous to the judgment, and (b) without awareness that the stereotyped information is a source of influence on judgment. (p. 278)

The findings by Todorov et al. (2004) provide another instance of social category applicability and extend its generality to the formation of implicit impressions.

Third, there is good evidence that characteristics of actors, including their social group memberships, influence the kinds of spontaneous inferences that occur. Dunning and Sherman (1997, Study 3) showed this with occupational stereotypes. Participants read sentences with ambiguous trait-implying behaviors, and rated how "readable" they were. Actor occupations were chosen that would disambiguate the behaviors in particular ways. Participants were then tested on their recognition memory for the sentences, with critical foils containing unambiguous variants of the sentences. False alarms were higher for sentence variants that were stereotype consistent than stereotype inconsistent. That is, stereotypes seemed to disambiguate the behaviors' meanings. These disambiguated meanings were stored in long-term memory and later led to accepting the unambiguous foils. Somewhat similarly, Delmas (1992; described in Uleman et al., 1996, p. 246) reported that actors' occupations facilitated occupational stereotype-related STIs, as seen in cued recall for the trait-implying sentences.

These studies, along with our speculations about the familiarity-validity heuristic described above, raise an interesting possible mechanism for the persistence of stereotypes. Once a stereotype is established, it can disambiguate ambiguous behaviors in a way that is consistent with the stereotype, serving to confirm it for particular individuals, as in Dunning and Sherman (1997). Then the mere repetition of that individual's face, with which the stereotype trait has become implicitly associated, should increase the perceived validity of a stereotyped perception of that person. All of this should occur without any explicit inference about the person or any awareness of the process. Of course, further research is needed to test this speculation directly.

The studies described above all concern the ways that particular stereotypes can affect implicit impressions of individuals, conditions that affect activation and/or binding of stereotype traits to actors. One other study concerns not particular stereotypes per se, but in-group/out-group effects. It illustrates

an additional way that social group membership may affect implicit impressions. Otten and Moskowitz (2000) examined whether membership in minimal groups would bias STIs. Participants read trait-implying sentences that described either in-group or out-group members—with membership established through a standard minimal group manipulation. Sentences were read in a recognition probe reaction time paradigm, in which participants had to judge whether a probe word appearing after each sentence had been explicitly present in the sentence. When the sentences implied the trait and described in-group members and the trait was positive, response times were longer. (In this paradigm, longer times to correctly decide that the implied trait had not been explicit in the sentence are evidence that the trait was inferred.) That is, in-group favoritism (but not out-group derogation) occurred, as evidenced by this spontaneous facilitation of positive in-group trait inferences. It is remarkable that this effect emerged without any explicit comparisons of in-group and out-group, on a measure that virtually assured participants' lack of awareness of what they were responding to and how.

So there is good research on the formation of implicit stereotypes and evidence that stereotypes and group membership affect implicit impressions in a variety of ways consistent with other research. More interestingly, these studies suggest that some of the methods used to study implicit impressions are also useful for increasing our understanding of the processes behind stereotyping and in-group/out-group phenomena.

Conclusion

The idea that we have implicit impressions of others—knowledge that we do not or cannot make explict, which nonetheless affects how we view and interact with others—is not new. It is captured in everyday speech about interpersonal "chemistry," "vibes," and intuition. Some of these intuitions come from social-cognitive transference, when the person resembles a significant other and, without realizing it, we assimilate the new person to the old one (e.g., chapter 16). And some of these intuitions come from past encounters with and inferences about that very person, but which we do not or cannot recall. These are the intuitions that this chapter is about.

STIs are implicit impressions because they are formed without intentions or awareness. There is now good evidence that they are inferences about people, not merely about behaviors. Carlston and Skowronski (1994; Carlston et al., 1995) showed this with their savings-in-relearning paradigm, and Todorov and Uleman (2002, 2003) showed it with their false recognition paradigm. Although STIs are occasionally associated with the "wrong" person—someone such as a communicator telling about the actor, as in Skowronski et al.'s (1998) STT—evidence from Todorov and Uleman (in press) suggests

that this is relatively rare. STT probably requires a much more vivid presentation of the communicator than the actor, such as a photo of the communicator but none of the actor.

The STT research provides the first rigorous demonstrations that incidental encounters with others leave impressions that persist long after the details of the encounter have faded from memory. (Hastie and Park's, 1986, now-classic distinction between online and memory-based impressions refers exclusively to intentionally formed impressions.) Participants in these studies spontaneously formed impressions from the behaviors they read about, and these impressions affected later behavior (relearning), even when memory for the behaviors was gone (Carlston & Skowronski, 1994; Carlston et al., 1995). Although Skowronski et al. (1998) emphasized that these implicit impressions were about the wrong person, the important points here are that they were formed without awareness (spontaneously), and they presumably became the sole basis for subsequent trait ratings because (on the evidence of Carlston et al.) the behaviors had been forgotten.

All of this suggests that whenever people encounter others they have seen before, their responses and reactions depend both on whatever they remember and on implicit impressions. Our adaptation of Jacoby's (1991) PDP provided a direct test and confirmed this suggestion. Right after an encounter, trait ratings of actors were based on both explicitly remembered (controlled) and implied (automatic) impressions. Two days later, only implicit impressions had any effect.

The spontaneity of STIs is a description of encoding processes: without intention or awareness. The implicitness of implicit impressions is a description of retrieval limitations: without explicit reference or memory. Presumably, implicit impressions can arise under many conditions. What makes them implicit is their memorial status at the moment. In PDP terms, they are the uncontrollable (automatic) effects of past encounters that affect current performance.

This formulation and set of paradigms suggest many questions for future research. What kinds of responses are particularly sensitive or insensitive to the impact of implicit impressions? What encoding and/or retrieval conditions affect this impact? The same questions can be, and have been, asked about explicit impressions (e.g., Wyer & Srull, 1989), but we cannot assume the answers are the same. In fact, the PDP model assumes A and C processes operate independently of each other. So reducing the impact of controlled processes need not increase the impact of implicit impressions.

More substantively, this formulation suggests directions for future research not only on implicit impressions, but also on STT, the bases for people's feelings that they know someone very well, stereotypes, in-group/out-group phenomena, and more.

Acknowledgments This research was conducted while Steven L. Blader was supported by the National Institute of Mental Health Grant T32-MH19890. The authors wish to thank Larry Jacoby and John Skowronski for their helpful comments on an earlier draft. We also thank Bianca Acevedo, Craig Herman, Lauren Pagano, Tarika Singh, Emily Thaden, and Kevin Young for their assistance in conducting some of the studies described here.

Notes

1. Part of the advantage of this false recognition measure over the cued recall measure used in prior research to detect actor-trait links may lie in the fact that false recognition is affected by both trait → actor links and actor → trait links (which also operate in savings in relearning). Prior research with trait-cued recall of actors tapped only the trait → actor links. It may be that actor → trait links are even stronger and formed more readily.

2. We did not test to see whether there was an interaction between delay and process type because it is not clear that the two process types are estimated on comparable scales.

References

Andersen, S. M., & Berenson, K. R. (2001). Perceiving, feeling, and wanting: Experiencing prior relationships in present-day interpersonal relations. In J. P. Forgas & K. D. Williams (Eds.), *The social mind: Cognitive and motivational aspects of interpersonal behavior* (pp. 231–256). New York: Cambridge University Press.

Anderson, C. A., Krull, D. S., & Weiner, B. (1996). Explanations: Processes and consequences. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 271–296). New York: Guilford.

Arkes, H. R., Hackett, C., & Boehm, L. (1989). The generality of the relation between familiarity and judged validity. *Journal of Behavioral Decision Making, 2*, 81–94.

Bacon, F. T. (1979). Credibility of repeated statements: Memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 241–252.

Banaji, M. R., Hardin, C., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology, 65*, 272–281.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition: Vol. 1. Basic processes* (pp. 1–40). Hillsdale, NJ: Erlbaum.

Bargh, J. A. (1999). The cognitive monster: The case 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.

Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin, 126*, 925–945.

Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of process in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General, 121*, 446–458.

Berk, M. S., & Andersen, S. M. (2000). The impact of past relationships on interpersonal behavior: Behavioral confirmation in the social-cognitive process of transference. *Journal of Personality and Social Psychology, 79*, 546–562.

Carlston, D. E., & Skowronski, J. J. (1994). Saving in the relearning of trait information as evidence for spontaneous inference generation. *Journal of Personality and Social Psychology, 66*, 840–856.

Carlston, D. E., Skowronski, J. J., & Sparks, C. (1995). Savings in relearning: II. On the formation of behavior-based trait associations and inferences. *Journal of Personality and Social Psychology, 69*, 420–436.

Chen, S., & Andersen, S. M. (1999). Relationships from the past in the present: Significant-other representations and transference in interpersonal life. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 31, pp. 123–190). San Diego, CA: Academic Press.

Claparède, E. (1951). Recognition and "me-ness." In D. Rapaport (Ed. & Trans.), *Organization and pathology of thought* (pp. 58–75). New York: Columbia University Press. (Original work published 1911)

Crawford, M. T., Sherman, S. J., & Hamilton, D. L. (2002). Perceived entitativity, stereotype formation, and the interchangeability of group members. *Journal of Personality and Social Psychology, 83*, 1076–1094.

Delmas, F. (1992, July). *Impact of target's category membership on spontaneous trait inference*. Poster presented at the XXV International Congress of Psychology, Brussels, Belgium.

Dunning, D., & Sherman, D. A. (1997). Stereotypes and tacit inference. *Journal of Personality and Social Psychology, 73*, 459–471.

Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review, 85*, 395–416.

Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 1, pp. 204–243). New York: Guilford.

Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 247–282). Mahwah, NJ: Erlbaum.

Ferguson, M. J., & Bargh, J. A. (2001). *Beyond response time: Consequences of automatic evaluative processes for the interpretation of subsequent stimuli*. Manuscript under review, New York University.

Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology, 58*, 203–210.

Fiedler, K., & Schenck, W. (2001). Spontaneous inferences from pictorially presented behaviors. *Personality and Social Psychology Bulletin, 27*, 1533–1546.

Gilbert, D. T. (1998). Ordinary personology. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (Vol. 2, 4th ed., pp. 89–150). Boston: McGraw-Hill.

Gill, M. J., Swann, W. B., Jr., & Silvera, D. H. (1998). On the genesis of confidence. *Journal of Personality and Social Psychology, 75*, 1101–1114.

Glassman, N. S., & Andersen, S. M. (1999). Activating transference without consciousness: Using significant-other representations to go beyond what is subliminally given. *Journal of Personality and Social Psychology, 77*, 1146–1162.

Greenwald, A. G., & Banaji, M. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4–27.

Greenwald, A. G., & Farnham, S. D. (2000). Using the Implicit Association Test to measure self-esteem and self-concept. *Journal of Personality and Social Psychology, 79*, 1022–1038.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74*, 1464–1480.

Hasher, L., Goldenstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior, 16*, 107–112.

Hassin, R., & Trope, Y. (2000). Facing faces: Studies on the cognitive aspects of physiognomy. *Journal of Personality and Social Psychology, 78*, 837–852.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review, 93*, 258–268.

Howard, P. (2000). *Racial knowledge effects on minorities: How stereotypic knowledge biases African-Americans' inferences*. Unpublished doctoral dissertation, New York University.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513–541.

Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Attention, awareness, and control. *Journal of Experimental Psychology: General, 122*, 139–154.

Jacoby, L. L., Yonelinas, A. P., & Jennings, J. M. (1997). The relation between conscious and unconscious (automatic) influences: A declaration of independence. In J. D. Cohen & J. W. Schooler (Eds.), *Scientific approaches to consciousness* (pp. 13–47). Mahwah, NJ: Erlbaum.

Lewicki, P. (1985). Nonconscious biasing effects of single instances on subsequent judgments. *Journal of Personality and Social Psychology, 48*, 563–574.

Lewicki, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 135–146.

Manis, M., Cornell, S. D., & Moore, J. C. (1974). Transmission of attitude-relevant information through a communication chain. *Journal of Personality and Social Psychology, 30*, 81–94.

Moskowitz, G. B., & Roman, R. J. (1992). Spontaneous trait inferences as self-generated primes: Implications for conscious social judgment. *Journal of Personality and Social Psychology, 62*, 728–738.

Newman, L. S., & Uleman, J. S. (1990). Assimilation and contrast effects in spontaneous trait inferences. *Personality and Social Psychology Bulletin, 16*, 224–240.

Otten, S., & Moskowitz, G. B. (2000). Evidence for implicit evaluative in-group bias: Affect-biased spontaneous trait inference in a minimal group paradigm. *Journal of Experimental Social Psychology, 36*, 77–89.

Schachter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 501–518.

Schachter, D. L., & Curran, T. (2000). Memory without remembering and remembering without memory: Implicit and false memories. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd edition, pp. 829–840). Cambridge, MA: MIT Press.

Skowronski, J. J., Carlston, D. E., Mae, L., & Crawford, M. T. (1998). Spontaneous trait transference: Communicators take on the qualities they describe in others. *Journal of Personality and Social Psychology, 74*, 837–848.

Stapel, D. A., Koomen, W., & van der Pligt, J. (1996). The referents of trait inferences: The impact of trait concepts versus actor-trait links on subsequent judgments. *Journal of Personality and Social Psychology, 70*, 437–450.

Todorov, A., & Uleman, J. S. (2002). Spontaneous trait inferences are bound to actors' faces: Evidence from a false recognition paradigm. *Journal of Personality and Social Psychology, 83*, 1051–1065.

Todorov, A., & Uleman, J. S. (2003). The automaticity of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology, 39*, 549–562.

Todorov, A., & Uleman, J. S. (in press). The person reference process in spontaneous trait inferences. *Journal of Personality and Social Psychology*.

Todorov, A., Gonzalez, C., Uleman, J. S., & Thaden, E. P. (2004). *A dissociation between spontaneous and intentional stereotyped trait inferences*. Unpublished manuscript, New York University.

Uleman, J. S., & Blader, S. (2001). *A process dissociation analysis of using spontaneous trait inferences*. Unpublished data, New York University.

Uleman, J. S., & Moskowitz, G. B. (1994). Unintended effects of goals on unintended inferences. *Journal of Personality and Social Psychology, 66*, 490–501.

Uleman, J. S., Newman, L. S., & Moskowitz, G. B. (1996). People as flexible interpreters: Evidence and issues from spontaneous trait inference. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 28, pp. 211–279). San Diego, CA: Academic Press.

Wegener, D. T., & Petty, R. E. (1997). The flexible correction model: The role of naïve theories of bias in bias correction. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 29, pp. 141–208). Mahwah, NJ: Erlbaum.

Wegner, D. M., & Bargh, J. A. (1998). Control and automaticity in social life. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (Vol. 1, 4th ed., pp. 446–496). Boston: McGraw-Hill.

Wegner, D. M., & Pennebaker, J. W. (Eds.). (1993). *Handbook of mental control*. Englewood Cliffs, NJ: Prentice Hall.

Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1235–1253.

Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology, 47*, 237–252.

Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology, 49*, 904–917.

Wyer, R. S., Budesheim, T. L., & Lambert, A. J. (1990). Cognitive representation of conversations about persons. *Journal of Personality and Social Psychology, 58*, 218–238.

Wyer, R. S., Budesheim, T. L., Lambert, A. J., & Swan, S. (1994). Person memory and judgment: Pragmatic influences on impressions formed in a social context. *Journal of Personality and Social Psychology, 66*, 254–267.

Wyer, R. S., Jr., & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Erlbaum.