Journal of Personality and Social Psychology
2004, Vol. 87, No. 4, 482–493

0022-3514/04/$12.00 DOI: 10.1037/0022-3514.87.4.482

The Person Reference Process in Spontaneous Trait Inferences

Alexander Todorov
Princeton University

James S. Uleman
New York University

Five studies examined whether spontaneous trait inferences uniquely reference the person who performed a trait-implying behavior. On each study trial in 5 studies, participants saw 2 faces and a behavioral sentence referring to one of them. Later, participants saw face–trait pairs and indicated whether they had seen the trait word in the sentence presented with the face. Participants falsely recognized implied traits more when these traits were paired with actors' faces than with control faces. This effect was replicated for a large set of faces (120), after a week delay between study and recognition test, when equal attention was paid to each face, and when the orientation of the face at recognition was different from the orientation at encoding.

People are remarkably good at spontaneously making inferences about others. Even when viewing 120 novel faces, each paired with a different behavior, people unintentionally and incidentally extracted specific trait information about each person (Todorov & Uleman, 2002). This occurs even when people's attention is divided and when the presentation of the information is very rapid (Todorov & Uleman, 2003). What is the nature of these spontaneous trait inferences "about" others? Does the trait become one of the logical predicates about the actor (as in "Jill is honest") or is it merely associated with the actor because of their co-occurrence ("Jill–honest")? Put another way, does the trait become part of the identity and representation of the person or is only the simple perceptual co-occurrence of faces and behaviors remembered?

Presumably, spontaneous social inferences occur under a greater variety of circumstances than intentional ones—certainly under a greater variety of goals and under a greater variety of observing and attentional conditions. You return from a party and remember that someone did something that was really gauche, but don't recall who it was. You notice an exchange between strangers in the elevator and later recall that one of them was overbearing, but you're not sure which one it was. While you're making dinner and reading the mail, you can't keep track of who is who in the BBC drama on TV; all those Brits look alike. To put these failures in terms of trait inferences, in some cases a trait is inferred but not associated with any particular actor. In other cases, it is associated with the wrong actor. There is good experimental evidence for both of these possibilities, as well as for getting it right spontaneously.

Uleman, Newman, and Moskowitz (1996) reviewed the cued-recall evidence on whether spontaneous trait inferences (STIs) are reliably linked or bound to the actors. Does reading that "the secretary solved the mystery half-way through the book" for a subsequent memory test make *clever* a reliable retrieval cue for *secretary*? The answer seemed to be no, except for special circumstances such as the reader's high Personal Need for Structure (Uleman et al., 1996, pp. 242–250).

On the other hand, Carlston and Skowronski (1994; Carlston, Skowronski, & Sparks, 1995) used a savings-in-relearning paradigm to show that STIs are reliably bound to actor representations, at least when actors are presented visually. In this paradigm, participants are presented with behavioral paragraphs and faces. Later, when asked to learn face–trait pairs, participants are faster to learn face–implied trait pairs than face–control trait pairs, indicating that the implied traits are inferred and linked to the actor representations. Furthermore, Skowronski, Carlston, Mae, and Crawford (1998; also Carlston et al., 1995, Experiment 4; Mae, Carlston, & Skowronski, 1999) showed that STIs could transfer to others, so that those who communicate about another person become associated with the trait implication of their own communication. The face–trait association was even obtained when participants were told that paragraphs and faces were randomly paired, and the genders of each did not match (Skowronski et al., 1998). These findings suggest that linking STIs to person representations could be a shallow associative process. In other words, such inferences might become implicitly associated with any face or object that happens to be present in the situation. Consistent with this suggestion, Brown and Bassili (2001) used the savings paradigm to show that STIs could even become associated with photographs of inanimate objects such as bananas and calculators.

Recently, Todorov and Uleman (2002) used a false recognition paradigm to provide evidence that STIs are linked to visual representations of the actor. Participants were presented with faces paired with single behaviors. Later in the experiment, participants were presented with face–trait pairs and asked to indicate whether they had seen the trait word in the sentence that had been presented with the face. Participants were more likely to falsely recognize implied traits paired with actors' faces than they were to recognize both implied traits paired with other faces presented earlier and novel traits paired with actors' faces. In a more recent set of studies, Todorov and Uleman (2003) argued that the process of binding STIs to actor representations is automatic. They replicated

Alexander Todorov, Department of Psychology, Princeton University; James S. Uleman, Department of Psychology, New York University.

We thank John Skowronski for his comments on drafts of this article and Fred DeLucia, Tarika Singh, Mike McLoughlin, and Anesu Mandisodza for their help with conducting experiments.

Correspondence concerning this article should be addressed to Alexander Todorov, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010. E-mail: atodorov@princeton.edu

the false recognition effects under conditions of rapid presentation, cognitive load, and shallow processing of the information. However, the question remains: Is this binding process anything more than a simple association?

The possibility that STIs are bound to any salient person representation is consistent with a general encoding specificity explanation of the false recognition and savings-in-relearning findings. The principle of encoding specificity states that "specific encoding operations performed on what is perceived determine what is stored, and what is stored determines what retrieval cues are effective in providing access to what is stored" (Tulving & Thomson, 1973, p. 369). Thus, it may be that faces or other objects (e.g., Brown & Bassili, 2001) are merely salient but incidental features of what is stored, providing a vivid context within which trait inferences are encoded. When this context is later provided, it facilitates the retrieval of trait inferences, resulting in a higher false recognition. The faces may function merely as any incidental context would, rather than showing that trait inferences are bound to actors in unique ways.

Uleman (1999) elaborated on this possibility when contrasting spontaneous with intentional impressions. He proposed that spontaneous impressions "are implicitly linked to actors by association, appropriately or not," whereas intentional impressions are "explicitly linked to the actors about whom one intends to form impressions" (p. 147). "Spontaneous impressions are implicitly associated with salient people, even if they are not the people the impression should be about" (Uleman, 1999, p. 155).

Yet, is this all there is to it? Do STIs become bound to any perceptually salient person representation? None of the false recognition studies presented more than one person at encoding. If both the actor and another person were present at encoding, would the STIs be preferentially bound to the actor? The current experiments were designed to explore this possibility by modifying the false recognition paradigm to include two faces (actor and control) and a behavior on each study trial. If STIs are simply bound to salient person representations, and the two faces are equally salient perceptually, then they should cue false recognition of implied traits equally. If, on the other hand, the STI process can (under the right conditions) include not only inferring a trait concept but also selectively attending to the relevant actor and linking the trait to that person representation spontaneously—that is, even though there is no impression formation goal—then the actor's face should cue false recognition of the implied trait more than the control face.

Modified False Recognition Paradigm

The false recognition paradigm consists of a study phase and a test phase. In the study phase, participants are presented with faces and behaviors implying traits (e.g., "Mary told the cashier that she got too much change"). In the test phase, participants are presented with the faces and traits (e.g., honest) and asked to decide whether they saw the trait in the sentence presented with the face. On some study trials, the trait is included in the sentence to make the recognition task believable (e.g., "Mary was honest and told the cashier that she got too much change"). In all previous experiments (Todorov & Uleman, 2002, 2003), in the study phase participants were presented with a single face and a behavior. In the present experiments, the study trials were modified to include two faces. In Experiments 1 to 3, the pair of faces was presented with a single trait-implying sentence that referred to one of the faces (actor). In Experiments 4 and 5, the pair of faces was presented with a pair of sentences, each sentence referring to one of the faces.

In the test phase, the implied trait was either paired with the actor's face or with the other (control) face that had been seen with the sentence. Instructions and practice trials made clear to participants that their task was to make a decision about a word from the sentence that had been presented with the face, whether actor or control. In other words, the recognition task did not require or even mention taking into account which person was referenced in the behavioral sentence.

If participants do not spontaneously form actor predicates in making spontaneous trait inferences, they should be as likely to falsely recognize implied traits in the context of control faces as in the context of actor faces. Similarly, the probability of correct recognition of explicitly presented traits should be the same in the two face conditions. In contrast, if participants do form actor predicates rather than mere associations in making trait inferences, they should show more false recognition of implied traits (and more correct recognition of explicit traits) in the context of actors' faces than control faces.

Experiments 1 and 2 were designed to test the hypothesis that STIs are more strongly associated with actors' faces than control faces. Experiment 3 explored the long-term effects of actor-linked STIs. It was modeled after Experiment 2, but participants were presented with the recognition test a week after the study phase. Experiments 4 and 5 addressed the possibility that the first three experiments' findings could be explained in terms of participants merely spending different amounts of time attending to actor and control faces. Experiment 5 tested whether the false recognition effects could be explained in terms of binding inferences to mere surface representations of faces rather than to representations of individuals.

Experiment 1

Participants were presented with 36 pairs of two faces and a behavior in the study phase of the experiment. Within each study trial, faces were matched on age, gender, and race. Each pair was presented for 10 s. Twelve of the behavioral sentences explicitly contained the implied trait word. In the test phase, participants were presented with 36 face–trait pairs and asked to decide whether they saw the trait word in the sentence presented with that face. If participants form actor predicates from implied traits rather than mere associations to them, they should be more likely to falsely recognize implied traits in the context of actor faces than control faces. Similarly, if they form predicates of explicitly presented traits at encoding, as part of the actor representations, they should correctly recognize such traits more in the context of actor faces.

Experiment 1 also manipulated participants' instructions. They had either memory instructions or impression formation instructions. Under impression instructions, they should intentionally infer traits about the actors. These participants were asked to think about the personality of the person described in each sentence, and they were told they would be asked about the person's character later in the experiment. If trait inferences are more likely to be associated with the person who performed the behavior under impression than memory instructions (e.g., Bassili & Smith, 1986), then effects should be stronger for impression participants. They should show higher false recognition of implied traits, and higher correct recognition of presented traits, when these are paired with the actors' rather than the control faces.

Method

Participants. Seventy-two undergraduates from the Department of Psychology at New York University (NYU) participated in the study for partial course credit. Participants were randomly assigned to four between-subjects conditions.

Stimulus material. Behavioral sentences were selected from those collected by Uleman and his students (Uleman, 1988). These were modified so that the pronouns were replaced by personal names. Thirty-six behavioral sentences were used in Experiment 1. These were pretested for their trait implications and used in Experiments 1 to 4 of Todorov and Uleman (2002).

Procedures. Participants worked individually in soundproof cubicles, and instructions were presented by a computer. The experiment consisted of a study phase and a test phase. Each study trial was presented for 10 s and the intertrial interval was 2 s. The order of the 36 trials was randomized for each participant. On each study trial, participants saw two faces, two names, and a behavioral sentence. The faces were in the upper left and right corners of the computer screen. The name of each person was below the person's face. The behavioral sentence was below the persons' names at the center of the screen. The sentence referred to one of the faces (the actor). For example, in one of the study trials the names were "Judith" and "Kim," and the behavioral sentence (implying *selfish*) was "Judith picked out the best chocolates before the guests arrived." Hence, Judith's face was the actor's face, and Kim's face was the control face. All pairs of faces were matched on race, gender, and age. The study phase started with a practice trial. In 12 of the 36 study trials, the behavioral sentence included the trait word (e.g., "Tom was so aggressive that he threatened to hit her unless she took back what she said.").

After the study phase, participants learned that in the second phase they would see faces from the first part of the experiment, each accompanied by a single word. They had to decide whether the word had been in the sentence presented with the person. If they thought it had been, they pressed the "OLD" key on the keyboard (the *M* key, labeled "OLD"). Otherwise they pressed the "NEW" key (the *X* key, labeled "NEW"). We asked participants to work as quickly as possible. To familiarize themselves with the task, they started with four practice trials. Before these trials, the pair of faces and the sentence from the practice trial in the study phase were presented for 5 s. Participants were reminded of the "OLD"/"NEW" decision. Then in two practice trials, each face (actor and control) was presented with a word that was part of the sentence. All participants pressed the "OLD" key and were given feedback. After these two trials, two more practice trials presented participants with each face and a word that was not part of the sentence. All participants pressed the "NEW" key. Participants were told that this was the correct decision, reminded to work as quickly as possible, and asked to continue with the experiment if everything was clear.

Each test stimulus was presented until the participant responded, and the delay between trials was 2 s. On each trial, a face was presented at the center of the screen with a trait below the face. The 24 pairs of faces that were presented with behavioral sentences in the study phase were randomly divided into two groups. Within each group of 12 pairs, each face was paired with the implied trait, thus creating 24 face–trait items (12 actor face–trait items and 12 control face–trait items) for each of the two groups. These were used to create two replication conditions, such that traits paired with actor faces in one replication were paired with control faces in the other replication, and vice versa. Thus, actor and control faces were counterbalanced in their pairing with implied traits across replication conditions. The same procedures were used for the 12 study trials in which the sentences included trait words. Within each replication, half the test trials presented the actor's face with the trait and half presented the control face with the trait.

At the beginning of the experiment, participants were given either memory or impression formation instructions. In the memory condition, they were told that the study was about how we remember information about other people and that their memory about each person would be tested later in the experiment. The type of memory test was not specified. In the impression condition, participants were told that the study was about how people form impressions of other people and how they determine what their personalities are like. Participants were asked to read the information and think about the character of the people in the sentences. They were told that in the second part of the study they would answer questions about these people's characters.

The overall design was a mixed 2 (instruction: impression vs. memory) $\times$ 2 (replication) $\times$ 2 (trait: presented vs. implied) $\times$ 2 (face: actor vs. control) analysis of variance (ANOVA) with the first two factors between-subjects and the last two within-subject. All analyses were performed at both the level of participants and the level of stimuli.

Results

As seen in Figure 1, recognition (both false and accurate) was higher when the trait word was paired with the actor's face (M = .62, SD = .17) than with the control face (M = .35, SD = .18), $F(1, 68) = 86.58$, $p < .001$. The rate of correct recognition of presented traits (M = .68, SD = .14) was higher than the rate of false recognition of implied traits (M = .30, SD = .19), $F(1, 68) = 235.91$, $p < .001$, reflecting some accuracy in memory. The overall Trait $\times$ Face interaction was significant, $F(1, 68) = 8.42$, $p < .005$, but more important, all four simple effects were also significant. False recognition of implied traits was higher when the traits were paired with the actor's face (M = .40, SD = .28) than when paired with the control face (M = .19, SD = .15), $t(71) = 7.38$, $p < .001$. Similarly, correct recognition of presented traits was higher when they were paired with the actor's face (M = .84, SD = .15) than when paired with the control face (M = .52, SD = .28), $t(71) = 8.00$, $p < .001$. For both actor and control faces, correct recognition of presented traits was higher than false recognition of implied traits, $ts(71) > 10.50$, $ps < .001$. None of the effects involving replication or instruction reached significance.

The analysis at the level of the stimuli was consistent with these results. However, in addition to the main effects of trait, $F(1, 34) = 187.95$, $p < .001$, and face, $F(1, 34) = 131.76$, $p < .001$, and their interaction, $F(1, 34) = 6.25$, $p < .017$, it also revealed a Face $\times$ Instruction interaction,[1] $F(1, 34) = 5.32$, $p < .027$. Whereas recognition rates for traits paired with control faces did not differ between the memory and impression conditions (M = .29, SD = .20 vs. M = .31, SD = .20, respectively), $t < 1$, the rates for traits paired with actor faces differed significantly, $t(35) = 2.55$, $p < .015$. Traits were more likely to be "recognized" as old (both accurately and falsely) under memory instructions (M = .57, SD = .24) than under impression instructions (M = .52, SD = .24).

Discussion

As expected, participants were more likely to falsely recognize implied traits when they were paired with actors' faces than with control faces. Unlike previous research using the false recognition

[1] Note that with stimuli as the unit of analysis, instruction is a within-subjects factor, rendering this analysis more sensitive to effects than when participants are the unit of analysis.

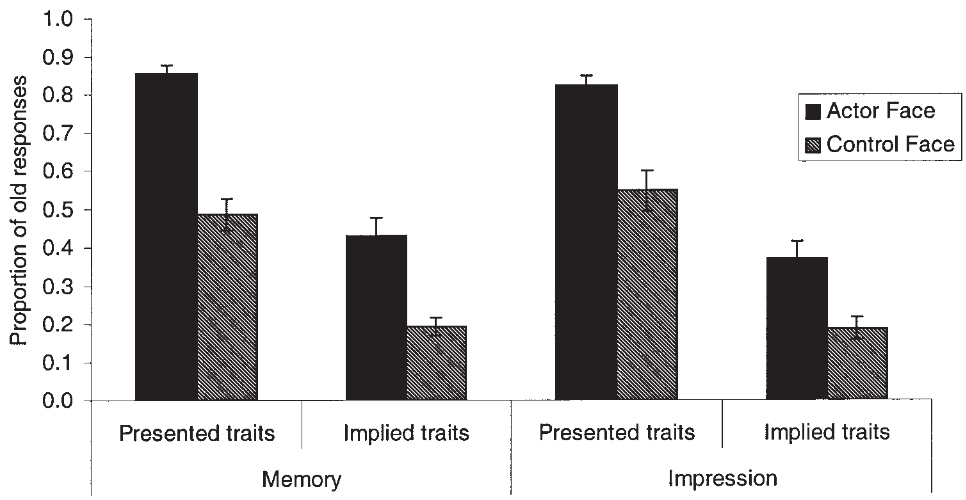

Figure 1. Proportion of old responses (correct for presented and false for implied traits) as a function of instructions at encoding and face at recognition. Experiment 1: Bars show standard errors.

paradigm (Todorov & Uleman, 2002, 2003), the actor and control faces were presented simultaneously with the behavioral sentence. The only difference between the two faces was that the trait-implying behavioral sentence referred to one of them. The higher false recognition in the context of the actor's face clearly shows that STIs are bound to actors. Further, this linkage does not depend on actors' perceptual salience, because both faces had equal perceptual salience.

Correct recognition of presented traits was similarly affected by the referential relation between faces and sentences. Participants were more likely to recognize presented traits in the context of the actor's face than in the context of the control face. This finding shows that both implicit and explicit attributes reference actors and that these attributes are encoded as part of the actor's representation. Interestingly, the pairing of traits with actors' faces did not increase overall recognition accuracy because, although pairing *presented* traits with actors' faces *increased* the correct response rates, pairing *implied* traits with actors' faces *decreased* correct response rates.

Experiment 1 found few differences between impression and memory participants. Although participants in the impression condition were asked to think about the actor, and in the memory condition they were not given any person-specific instruction, false recognition of implied traits and correct recognition of presented traits were affected in the same way by the face manipulation. If anything, the Face × Instruction interaction at the level of stimuli suggested that the association of traits with actors was stronger for memory participants.

This small effect of instructions is consistent with previous findings from paradigms using faces as critical stimuli to measure associations between STIs and actors. Carlston and Skowronski (1994) did not find differences between impression instructions and other instructions such as familiarization with the material and memorization, across five experiments using the savings paradigm (see also Carlston et al., 1995, Experiment 1). Todorov and Uleman (2002) argued that binding STIs to faces is a fairly robust phenomenon, relatively independent of processing instructions and processing time. Hence, it is difficult to demonstrate differences in processing as a result of instructions unless one manipulates other relevant factors that can show dissociative effects on impression and memory instructions. In one such experiment, Todorov and Uleman (2002, Experiment 4) showed that doubling the presentation time of face–behavior pairs reduced the probability of false recognition under memory instructions but had no effect under impression instructions. The current experiment did not include such factors as presentation time that could have had dissociative effects on memory and impression participants. So the findings are inconclusive with respect to other possible differences between memory and impression instructions.

Experiment 2

Experiment 1 used a counterbalancing scheme in which implied traits appeared later either with actor or control faces. However, the associations between particular faces and behaviors were not counterbalanced. Although we attempted to match actor and control faces on similarity, it is possible that some pairings of actor faces and behaviors were more likely to elicit trait inferences than pairing control faces with these same behaviors. People are willing to "read" specific traits into faces (e.g., Hassin & Trope, 2000). The actors' faces might have "conveyed" information that was more consistent with the implied trait than the information conveyed by the control faces. Experiment 2 was designed to replicate Experiment 1 and to rule out an explanation of its effects in terms of specific face–behavior combinations.

In this experiment, the association between particular faces and behaviors was manipulated. If a face served as the actor's face in one of the experimental conditions, it served as a control face in the other experimental condition. If the effects observed in Experiment 1 are due to the links formed at encoding between faces and behaviors, as we claim, rather than to specific combinations of faces and behaviors, participants should still be more likely to falsely recognize implied traits and correctly recognize presented traits in the context of the actor's face, even when the same faces serve as actors and as controls across the trials.

Experiment 2 also increased the number of faces and behaviors. Participants were presented with 120 faces and 60 behaviors (2 faces and a behavior per trial). Todorov and Uleman (2002, Experiments 5 & 6) showed that even when participants were presented with 120 faces, each paired with a different behavior, they were more likely to falsely recognize implied traits in the context of the actor's face than in the context of a face randomly paired with the trait. Thus, under conditions of information overload, participants were able to extract information specific to individuals. The present task is even more difficult because the actor and control faces are simultaneously presented with the behavior. Replicating the false recognition effect under these conditions would provide strong evidence for STIs' links to actors.

Method

Participants. Twenty-nine undergraduate students from the Department of Psychology at NYU participated in the study for partial course credit. Participants were randomly assigned to two between-subjects conditions.

Procedures. All participants received memory instructions. The procedures were the same as in Experiment 1 with two exceptions. The set of faces–behavior pairs was increased from 36 to 60, and the counterbalancing scheme was different. The 60 behaviors were randomly assigned to 60 new pairs of faces. Then two replications conditions were created. Unlike Experiment 1, in which counterbalancing was in the test phase, in this experiment it was in the study phase. Specifically, within each pair of faces, if a face was the actor in the behavioral sentence in one of the replications conditions, it was the control face in the other. For example, if "Judith picked out the best chocolates before the guests arrived" in one of the replications, then in the other replication, "Kim picked out the best chocolates before the guests arrived." Of the 60 sentences, 20 contained the implied trait word.

In the test phase, a face was randomly selected from each pair of faces and paired with the implied trait word. There were two constraints on this process. First, the implied trait was paired with the actor's face on half of the test trials and with the control face on the other half. Second, the implied trait was paired with the face presented on the right side of the screen in the study phase on half the trials and with the face presented on the left on the other half. The overall design was a mixed 2 (replication) × 2 (trait: presented vs. implied) × 2 (face: actor vs. control) ANOVA with the first factor between-subjects and the latter two within-subject.

Results

As shown in Figure 2, participants were more likely to report seeing the word for both implied and presented traits when the trait was paired with the actor's face (M = .62, SD = .12) than when paired with the control face (M = .46, SD = .18), $F(1, 27) = 16.09$, $p < .001$. Participants were more likely to correctly recognize presented traits (M = .67, SD = .12) than to falsely recognize implied traits (M = .42, SD = .14), $F(1, 27) = 63.52$, $p < .001$. The analysis also revealed a main effect of replication, $F(1, 27) = 6.43$, $p < .017$, indicating that participants were more likely to report seeing a trait in one of the replications than in the other. However, replication did not interact with any of the other factors.

Consistent with the above analyses, analyses at the level of stimuli showed the same pattern, revealing a significant main effect of face, $F(1, 58) = 38.21$, $p < .001$, and of trait, $F(1, 58) = 51.12$, $p < .001$.

Discussion

Participants were more likely to falsely recognize implied traits in the context of the actor's face than in the context of the control face. They were also more likely to correctly recognize presented traits in the context of the actor's face. Thus, this experiment replicated Experiment 1 with a larger set of faces and behaviors and with a counterbalancing scheme that ruled out an explanation in terms of particular face–behavior pairings. The findings strongly suggest again that STIs refer to the actors rather than to others who are simultaneously presented with the actors.

Experiment 3

In both Experiments 1 and 2, the recognition test immediately followed the study phase of the experiment. However, it is possible that the effect of binding STIs to actors' faces is short-lived and easily dissipates over time, especially with a large number of faces and behaviors.

Experiment 3 was designed to explore the long-term effects of binding STIs to actors' faces. It was modeled after Experiment 2, but the recognition test was presented a week after the study phase.

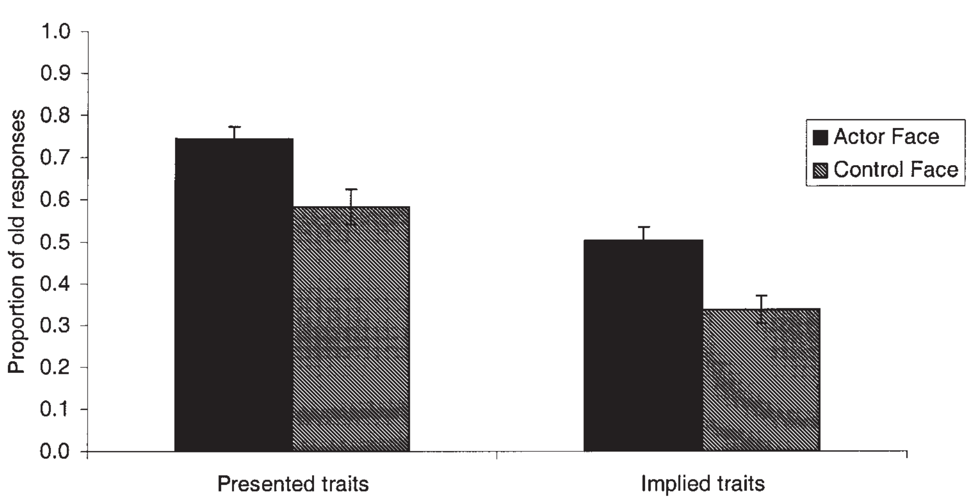

Figure 2. Proportion of old responses (correct for presented and false for implied traits) as a function of face at recognition. Experiment 2.

If the association between traits and actors is preserved in memory, participants should still be more likely to falsely recognize implied traits and to correctly recognize presented traits in the context of the actor's face than in the context of the control face after a week's delay.

Method

Participants. Thirty-two undergraduates from the Department of Psychology at NYU participated in the study for partial course credit. Participants were randomly assigned to two between-subjects conditions.

Procedures. The stimuli and the procedures were the same as in Experiment 2. The only difference was that the recognition test occurred a week after the study phase. The overall design was a mixed 2 (replication) $\times$ 2 (trait: presented vs. implied) $\times$ 2 (face: actor vs. control) ANOVA with the first factor between-subjects and the last two within-subject.

Results

The only significant effect was the effect of face, $F(1, 30) = 5.50$, $p < .026$. As shown in Figure 3, participants were more likely to report seeing the trait word for both implied and presented traits when the trait was paired with the actor's face ($M = .52$, $SD = .11$) than when paired with the control face ($M = .46$, $SD = .13$). This was the only significant effect at the level of stimuli too, $F(1, 58) = 5.78$, $p < .019$.

Discussion

Even after a week, participants were more likely to falsely recognize implied traits and to correctly recognize presented traits in the context of actor faces than control faces. Although these effects were not as strong as when the recognition test occurred immediately, they replicated reliably. In both Experiments 1 and 2, the probability of correct recognition of presented traits was much higher than the probability of false recognition of implied traits. However, after a week the difference between correct recognition and false recognition virtually disappeared. The respective proportions were identical ($M = .49$, $SD = .15$ vs. $M = .49$, $SD = .09$). That is, after a week the recognition performance was at a chance level. Nevertheless, the effect of face remained reliable.

These findings suggest that the implications of implied and presented traits are the same in the long term. Whether the trait was explicitly presented or was only inferred, it did not matter after a substantial delay. After a week, the difference between explicitly presented person attributes and inferred person attributes disappeared, but the link between attributes and persons remained.

Experiment 4

Across all three experiments, participants were more likely to falsely recognize implied traits in the context of actor faces than control faces. This finding strongly suggests that STIs were specifically bound to the actor, rather than being associated with any face that happened to be presented simultaneously with the behavior. This suggests the operation of processes other than (or in addition to) mere associations.

However, there may be a simpler explanation of these results, in terms of spending different amounts of time attending to actors' and control faces while the STI is activated. According to this explanation, the sentence points to a specific face and participants spend more time attending to this face during the 10-s trial. For example, a participant might look 7 s at the actor's face and only 3 s at the control face. In that way, the actor's face would provide an encoding context for the STI for a longer time than the control face, produce stronger associations, and thus account for the higher rates of both false and correct recognition.

Of course, we argue that spending more time attending to the actor's face is part of the STI process itself, a spontaneous allocation of attentional resources. However, we also think that mere time attending is not the whole explanation for the difference between actor and control faces. Therefore, the objective of Experiment 4 was to rule out this explanation. To control for the amount of time spent attending to each face, each study trial was modified to include two behaviors and two faces. The two behavioral sentences implied different traits, had different persons as topics, and thus referred to different faces. For example, if the first

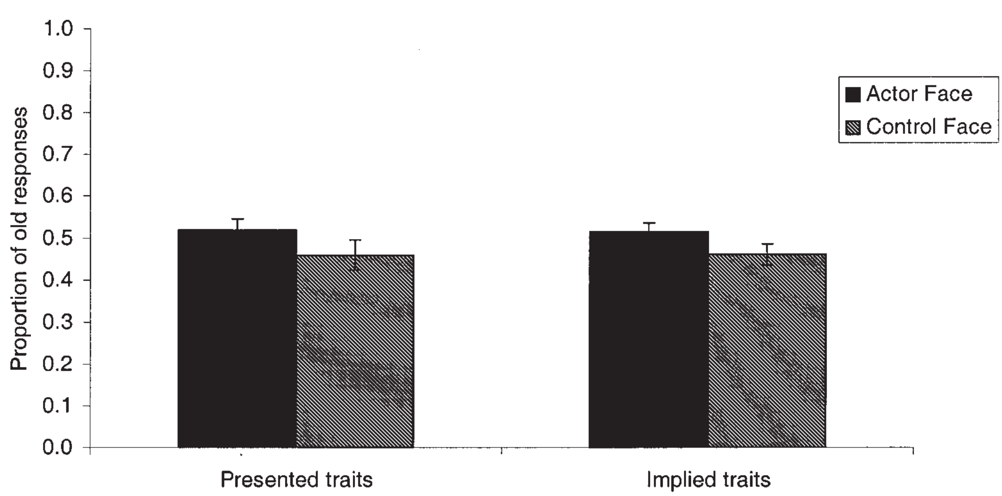

Figure 3. Proportion of old responses (correct for presented and false for implied traits) as a function of face at recognition, 1 week after study session. Experiment 3.

sentence referred to the face on the right, the second sentence referred to the face on the left. A critical feature of this design is that participants have to attend to both faces for the same amount of time. Both faces serve simultaneously as actor and control faces, because this distinction only occurs with respect to the implied trait presented at the recognition test. That is, depending on which implied trait is presented, the same face can function as either an actor's face or a control face. Although participants can switch their attention between the faces when reading the behavioral sentences, they have to allocate the same amount of time and attention to each face. If the observed recognition effects in the first three experiments are not entirely due to differential attention to actor and control faces (and we hypothesize they are not), then participants should be more likely to falsely recognize implied traits and to correctly recognize presented traits in the context of actor faces than control faces, in spite of having attended equally to both.

Method

Participants. Twenty-six undergraduates from the Department of Psychology at NYU participated in the study for partial course credit. Participants were randomly assigned to two between-subjects conditions.

Procedures. The general procedures were the same as in the previous experiments. All participants were given memory instructions. Each study trial was presented for 10 s, and the order of trials was randomized for each participant. The order of the test trials was also randomized. The intertrial interval was 2 s for both study and test trials.

Most important, the study trials in this experiment were modified to include two behaviors implying different traits. Thus, each trial consisted of two faces presented in the right and left corners of the screen, two names below the faces, and two behavioral sentences below the names and centered on the screen. If the first sentence referred to the person on the right, then the other sentence referred to the person on the left, and vice versa. For example, one of the study trials presented the faces of Bill and Ted with the following sentences: "Bill thought he didn't deserve their award and praise," implying "modest," and "Ted carried the old woman's groceries across the street," implying "helpful." Participants were presented with 36 study trials. In 12 of these, the sentences contained the implied trait.

In this design, there are four possible test trials for each study trial. Each face can be paired either with the implied trait in the sentence about the face (actor pairing) or with the implied trait in the sentence about the other face (control pairing). In the example above, the two actor pairings are Bill–modest and Ted–helpful, and the two control pairings are Bill–helpful and Ted–modest. To avoid presenting the same trait twice or the same face twice, if one of the faces was paired with the implied trait about it (e.g., Bill–modest), the other face was also paired with its implied trait (e.g., Ted–helpful). For half of the test trials, the faces were paired with the implied traits about them (actor pairing), and for half they were paired with the implied traits about the other face (control pairing). Two replications conditions were created so that actor and control pairing were counterbalanced across participants. For example, if Bill and Ted were paired with their implied traits (Bill–modest and Ted–helpful) in one of the replication conditions, they were paired with the control traits in the other condition (Bill–helpful and Ted–modest). The overall design was a mixed 2 (replication) $\times$ 2 (trait: presented vs. implied) $\times$ 2 (face: actor vs. control) ANOVA with the first factor between-subjects and the latter two within-subject.

Results

As shown in Figure 4, participants were more likely to report seeing a trait when it was paired with the actor's face (M = .60, SD = .11) than with the control face (M = .49, SD = .10), $F(1, 24) = 13.20$, $p < .001$. They were also more likely to correctly recognize presented traits (M = .65, SD = .09) than to falsely recognize implied traits (M = .45, SD = .09), $F(1, 24) = 122.64$, $p < .001$. Analyses at the level of stimuli showed the same pattern: $F(1, 70) = 10.93$, $p < .001$, for the main effect of face, and $F(1, 70) = 37.28$, $p < .001$, for the main effect of type of trait. There were no other significant effects.

Discussion

In contrast to the first three experiments, in which the actor and control faces were clearly designated in each study trial, each trial in Experiment 4 presented behavioral information about each face. Thus, the distinction between actor and control faces did not occur until the recognition test and depended only on which trait was presented with a face. Yet, as in the first three experiments, participants were more likely to falsely recognize implied traits and to correctly recognize presented traits in the context of actor faces than control faces. These findings strongly suggest that the prior findings cannot be explained in terms of differential attending to actor and control faces for different amounts of time. Instead, the findings must be accounted for in terms of participants' spontaneous processing during that time.

Experiment 5

The first four experiments provided clear evidence that STIs were bound to actors' faces. However, the more important question is whether such inferences are bound to a more abstract representation of the individual, a representation that should be independent of the particular view of the person's face. It is possible that the observed recognition effects were due to binding of inferences to particular perceptual configurations of faces rather than to a representation of the person. In fact, all previous experiments in both the savings and the false recognition paradigm have used the same person photographs at study and at test.

To test whether STIs are encoded as part of the representation of the individual, we manipulated the orientation of the face at study and at test by using two different pictures of the same individual from different angles. Because this manipulation changes the perceptual context away from that at encoding, we expected that it would decrease the overall rate of both correct recognition of presented traits and false recognition of implied traits. More important, if STIs are encoded as part of an abstract representation of the individual, participants should be more likely to falsely recognize implied traits and correctly recognize presented traits in the context of the actor's face than in the context of the control face even when the pictures of these individuals are different at test. That is, STIs should be bound to a view-independent representation of the actor. Finally as in Experiment 4, the distinction between actor and control faces does not occur until the recognition test, because the faces serve as both actor and control simultaneously.

Method

Participants. Fifty-eight undergraduate students at Princeton University participated to fulfill research requirements for introductory level

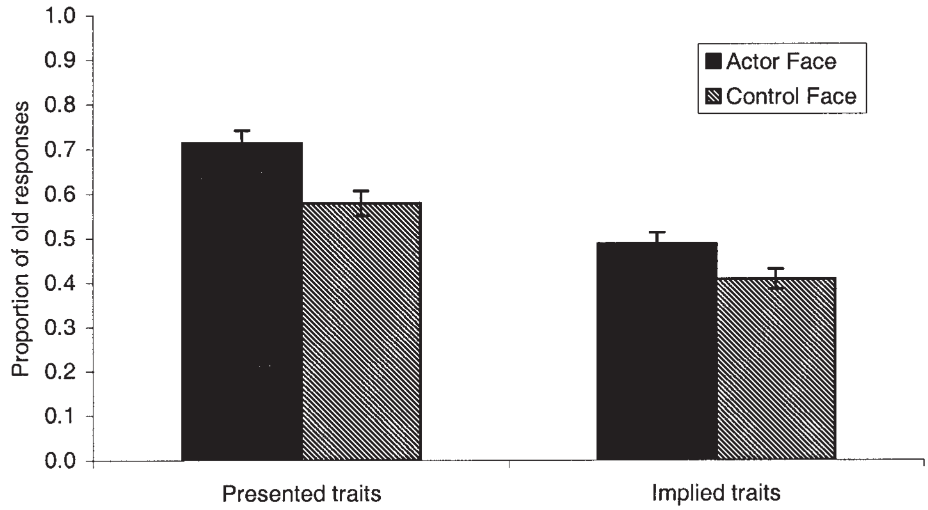

Figure 4. Proportion of old responses (correct for presented and false for implied traits) as a function of face at recognition. Study sentences referenced both faces. Experiment 4.

psychology courses. Participants were randomly assigned to four between-subjects conditions.

Stimuli. For this experiment, we used a set of novel faces taken from a previously established database (Lundqvist, Flykt, & Öhman, 1998). The database consists of pictures of 70 amateur actors, 35 women and 35 men, between 20 and 30 years of age. In the pictures, all actors wore gray T-shirts and there were no beards, mustaches, earrings or eyeglasses, or visible make-up. For the experiment, we selected pictures of 64 individuals with neutral expressions (32 men and 32 women) in two different face orientations: full forward and half right profile.

Procedures. The general procedures were the same as in Experiment 4. Sixty-four behaviors were randomly assigned to the 64 faces. In the study phase, participants were presented with 32 study trials. Each study trial consisted of 2 faces matched on gender and 2 behaviors matched on valence. In 8 of the trials, the sentences contained the implied trait. On all study trials, the faces were in the full forward orientation.

The major difference from Experiment 4 was that in the test phase, the faces were presented either in the full forward orientation as in the study phase or in half right profile. Four replications conditions were created so that actor and control pairing in the two face orientations were counterbalanced across participants. Specifically, the 64 faces were randomly divided into four groups of 16 faces (8 men and 8 women), and each group of faces was presented in each of the four possible combinations (actor pairing–straight face, control pairing–straight face, actor pairing–half right profile face, control pairing–half right profile face) across participants at test. The overall design was a mixed 4 (replication) $\times$ 2 (trait: presented vs. implied) $\times$ 2 (face: actor vs. control) $\times$ 2 (face orientation: same as in study phase vs. different) ANOVA with the first factor between-subjects and the latter three within-subjects.

Results

Replicating Experiment 4 and as shown in Figure 5, participants were more likely to report seeing a trait when it was paired with the actor's face ($M = .51$, $SD = .15$) than with the control face ($M = .42$, $SD = .14$), $F(1, 54) = 23.00$, $p < .001$. They were also more likely to report seeing a trait when the orientation of the face was the same as in the study phase ($M = .50$, $SD = .15$) than when the orientation was different ($M = .43$, $SD = .15$), $F(1, 54) = 13.71$, $p < .001$. This effect was qualified by an unexpected interaction of replication and orientation, $F(3, 54) = 3.61$, $p < .019$, indicating that in one of the four replication conditions the rate of recognition for faces presented in the same orientation was not higher than the rate for faces presented in a different orientation. As in the previous experiments, participants were more likely to correctly recognize presented traits ($M = .57$, $SD = .17$) than to falsely recognize implied traits ($M = .37$, $SD = .15$), $F(1, 54) = 76.16$, $p < .001$. No other effects reached significance. However, the interaction of face and orientation approached significance, $F(1, 54) = 3.61$, $p < .063$, suggesting that the effect of actors' faces relative to control faces was weaker when the face orientation was changed at test. An analysis of simple effects showed that the higher reporting of traits in the context of the actor's face was reliable even when the orientation of the face was different, $t(57) = 2.19$, $p < .033$.

Analyses at the level of stimuli showed the same pattern: $F(1, 62) = 8.65$, $p < .005$, for the main effect of face, $F(1, 62) = 13.97$, $p < .001$, for the main effect of face orientation, and $F(1, 62) = 64.98$, $p < .001$, for the main effect of type of trait. There were no other significant effects.

Discussion

Replicating Experiment 4, participants were more likely to report seeing a trait in the context of the actor's face than in the context of the control face. It should be emphasized that in both experiments, each face served as both actor and control and, hence, the rate of recognition was completely determined by the trait pairing at test. Most important, the effect of actor versus control faces was obtained even when the faces at test were presented from a different angle. This finding strongly suggests that the trait inference was bound to the representation of the individual rather than to a particular perceptual face configuration. Predictably, the face orientation manipulation reduced the overall rate of both correct recognition of presented traits and false recognition of implied traits, but it did not eliminate the actor effect.

The findings of Experiments 4 and 5 rule out an explanation of the prior findings in terms of ignoring the control face during the study time. In the last two experiments, participants had to pay

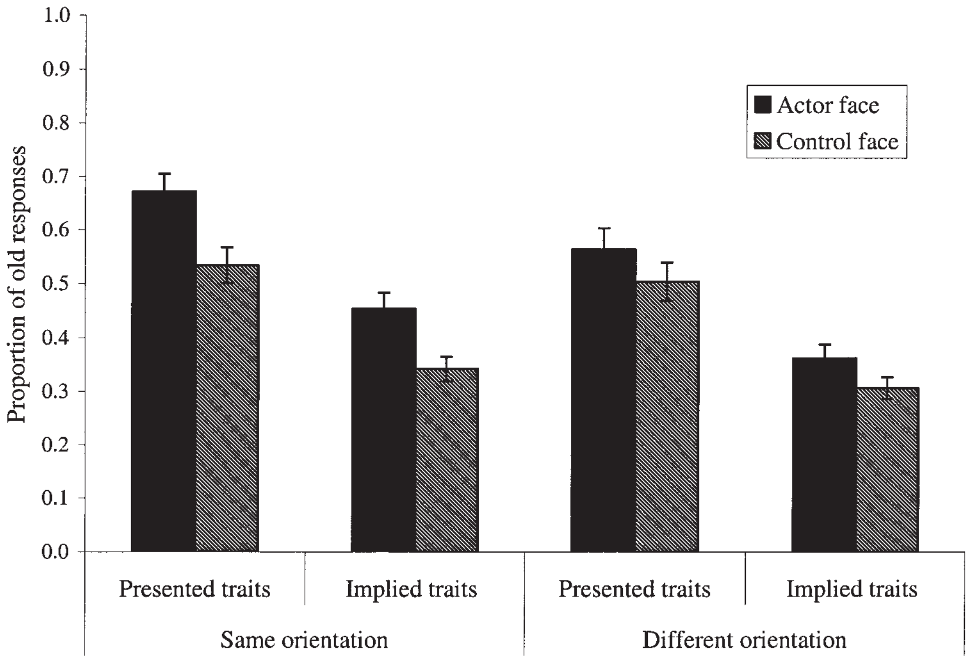

Figure 5. Proportion of old responses (correct for presented and false for implied traits) as a function of face at recognition and orientation of face at recognition. Study sentences referenced both faces. Experiment 5.

attention to both faces, because both faces served as actors on each trial. It is quite plausible that participants divided their attention between the two faces while reading the behavior information, but the design of the experiments precludes the possibility of preferential attention to one of the faces. At the same time, it seems that such preferential attention did play some role in the first two experiments. The effect sizes in terms of Pearson's correlation for the false recognition of implied traits were .66 and .62 in Experiments 1 and 2, respectively. (Experiment 3 is not relevant because there was a 1-week delay between study and test.) The corresponding effect sizes for Experiments 4 and 5 were .44 and .47, respectively (trials where the face orientation was the same at study and test). A contrast comparison of these effect sizes (Rosenthal & Rubin, 1982) approached significance, $z = 1.59$, $p < .06$ (one-tailed), suggesting that the effects in Experiments 1 and 2 were partially driven by preferential attention to actors' faces (e.g., Fiske, 1980). Nevertheless, it is clear from the data that binding STIs to actors' representations is also driven by other mechanisms that are not explicable simply in terms of amount of attention.

General Discussion

If spontaneous trait inferences are bound specifically to the representation of the actor and not to the encoding context in general, participants should be more likely to falsely recognize implied traits paired with the actors' faces than the same traits paired with a control face that had been presented simultaneously. This prediction was confirmed in five experiments. Participants were more likely to falsely recognize implied traits in the context of the actor than the control faces, even a week later (Experiment 3), when the task equated the amount of time spent attending to each face (Experiments 4 and 5), and when the face of the actor was shown from a different angle at the recognition test than at encoding (Experiment 5). These five experiments provide critical evidence that STIs are about the actor—in the sense of being properties or logical predicates of the actors—rather than being simply associated with any salient person.

The correct recognition of explicitly presented traits was affected in the same way as the false recognition of implied traits. Participants were more likely to recognize the traits when they were paired with the actor than with control faces. As a result, the pairing of the traits with actors increased both correct recognition of presented traits and false recognition of implied traits. These recognition findings suggest that both implicit and explicit person attributes are encoded as part of the actor representation.

Although after a week correct recognition was the same as false recognition, the effect of the face remained reliable. Participants were more likely to associate the trait, independently of its implied or explicit nature, with the actor. These findings suggest that the implications of inferred and explicit traits for the actor representation are the same in the long term. Further, although correct recognition of presented traits was substantially higher than false recognition of implied traits in all four experiments without a delay, the effect size for associating traits with actors rather than controls was equivalent. Across the four experiments it was .57 (Pearson's r). That was also the effect size for the higher rate of correct recognition of presented traits in the context of the actors' faces.[2] Thus the strength of the association between actors and traits was the same for inferred and explicit traits.

[2] The effect sizes for Experiment 5 were computed on the trials in which the face orientation was the same at study and at test, because these trials were equivalent to the trials in the other experiments.

Contributions of the Modified False Recognition Paradigm

The findings of the five experiments are consistent with the findings of the experiments reported in Todorov and Uleman (2002, 2003). On each study trial in the latter experiments, participants were presented with a face and a behavior rather than with two faces and a behavior (or two behaviors), as in the current experiments. In this single face paradigm, participants were more likely to falsely recognize implied traits in the context of the actor's face than in the context of a control face, which was presented earlier in the experiment, or novel traits in the context of the actor's face.

The modified false recognition paradigm makes two new and related contributions: It provides a better control for the measurement of actor-linked STIs and it provides more specific evidence on the nature of the mechanism underlying binding inferences to person representations. In this paradigm, the control face is presented simultaneously with the actor's face and the behavioral information. With this presentation, the inferred trait is necessarily encoded in the context of both the actor and the control face, potentially reducing or eliminating the difference in effects of actor and control faces on trait recognition. The recognition task does not require discrimination between the faces. Participants are simply asked to make a decision about information presented with the face. The correct response is the same whether the test face is the actor's face or the control face. Nevertheless, participants discriminated between these faces for both inferred and explicit traits, so that actor and control faces had sharply different effects.

Although the false recognition effect seems more difficult to obtain in the modified false recognition paradigm, comparison of the effect sizes for the two paradigms shows only a small difference. The average effect size across the six experiments reported in Todorov and Uleman (2002) was .66 (Pearson's r), whereas the average effect size across the four experiments reported here was .57. (Experiment 3 with a week delay was excluded, because none of the other experiments included a delay.) Comparisons between different paradigms can be misleading, but this similarity of effect sizes does suggest a powerful mechanism binding inferences to particular representations. The current findings show that STIs are not simply bound to any face that happens to accompany the trait-implying information. STIs are about particular persons and are bound to representations of those persons.

General Associative Mechanisms Versus Predicate-Producing Inferences

The findings from the modified false recognition paradigm show that STIs reference the actor. At the same time, spontaneous trait transference (STT) findings from the savings paradigm show that STIs are associated with any face that is presented with the behavioral information (Mae et al., 1999; Skowronski et al., 1998). These two sets of findings may seem to be inconsistent, but they are not. In the first study that demonstrated the transference effect, Carlston et al. (1995) noted that "the [person–trait] associations formed when the behaviors pertain directly to the target are stronger" (Experiment 4, p. 430). As reported below, a meta-analysis of subsequent STT studies confirms this hypothesis.

In STT studies, participants are told that the person in the photograph is telling information about a different person, the person described in the trait-implying paragraph. In three of these experiments, the study and the learning tasks were in the same experimental session (Carlston et al., 1995, Experiment 4; Mae et al., 1999, Experiments 1 & 2). The average effect size for these experiments in terms of Pearson's r was .36. In contrast, the effect size was .75 for two STI experiments (Carlston & Skowronski, 1994, Experiment 1; Carlston et al., 1995, Experiment 3) in which the descriptions were about, rather than told by, the person in the photograph, and the study and learning tasks were in the same session. A contrast test of these two sets of studies showed a highly significant difference ($z = 5.26$, $p < .001$). In fact, the effect size for the three STI experiments (Carlston & Skowronski, 1994, Experiments 2 & 4; Carlston et al., 1995, Experiment 1) that introduced a 2-day delay between the study and learning tasks was as large ($r = .74$) and also significantly different from the effect size of the STT experiments ($z = 4.13$, $p < .001$).[3] Brown and Bassili's (2001) findings are consistent with these differences in effect sizes.

Thus, findings from both the savings and the false recognition paradigms clearly show that STIs are more strongly associated with the actors' faces than other faces. But the findings from the savings paradigm also show that spontaneous inferences can become associated with faces besides the actor's face. This is not inconsistent with the false recognition findings that STIs are more strongly linked to actors' than other faces when both are presented simultaneously. The other faces presented simultaneously in these modified false recognition studies served as our controls. It remains to be seen whether, relative to controls such as faces presented on different trials, there might also be STTs, that is, associations between inferred traits and simultaneously presented nonactor faces. It is plausible that spontaneously inferred traits can be more strongly associated with simultaneously presented nonactor faces than with faces seen on other trials in the false recognition paradigm. This hypothesis remains for future research.

The current false recognition findings certainly do not rule out a shallow associative process that binds unintentional inferences to objects other than their referents, such as messengers delivering trait-implying messages about others. That is, inferences may get bound to any objects presented simultaneously with the information triggering the inferences, including Brown and Bassili's (2001) suspicious bananas. However, the larger general issue is whether the same shallow process can explain differences in the strength of bonds between unintended inferences and various objects. For this to succeed, as yet unidentified factors that determine the strength of these bonds must be discovered. Alterna-

[3] The formula and the rationale for the contrast tests can be found in Rosenthal and Rubin (1982). The STT studies and the studies in which the description referred to the person in the photo were assigned equal contrast weights but different signs. There were actually four STT experiments that used the savings task, but the fourth experiment (Skowronski et al., 1998, Experiment 1) had two delay conditions (no delay and a 2-day delay) and the relevant statistics for the separate conditions were not reported. However, the effect size across the two conditions (.35) was quite comparable to the effect size for the no delay experiments (.36). Similarly, there were eight savings experiments in which the description pertained to the person, but we used only five in the comparison because either we were unable to derive relevant statistics because of multiple conditions (Carlston & Skowronski, 1994, Experiments 3 & 5) or the inferences were clearly intentional (Carlston et al., 1995, Experiment 2).

tively, one could posit a spontaneous rule-based reference mechanism that operates automatically and in addition to general associative binding mechanisms. This would account for the consistent superiority of actor faces over control faces in producing false recognition in these studies.

Then the interesting question becomes, does the inferred trait reference an abstract person representation or merely the specific perceptual representation of that person presented at encoding? Our last study, showing that the effect occurs even when the face's orientation changes from study to test, supports the idea that the trait refers to an abstract representation. That is, not only are trait concepts spontaneously inferred from behavior, but they are also spontaneously linked to actor representations that extend beyond mere surface perceptual features.

An important question for future research is whether there is something unique about spontaneous person inferences relative to spontaneous inferences about the properties of inanimate objects. Would reading that "the vase shattered into hundreds of pieces when it fell from the coffee table" prompt an inference of "fragile" and link that inference to that vase more than to others pictured at the time, even when it was seen from a different perspective? Much research on "theory of mind" shows that inferences that people make about animate agents differ from those they make about inanimate objects (e.g., Lillard & Skibbe, in press; Malle, in press; Wellman, 1990). Some of this may be cultural, that is, acquired and not based on specialized innate mechanisms. However, there is also evidence from primatology that people have uniquely evolved to make complex inferences about other agents' mental states (e.g., Tomasello, 1999). There also seems to be good neuroscience evidence for specific brain regions devoted to processing information about animate agents rather than inanimate objects (e.g., Frith & Frith, 1999; Gallagher & Frith, 2003). All of this suggests that inferences about agents are privileged in some way, over associations to objects, and involve cognitive processes more complex than simple associations.

This issue is also related to the debate on face recognition that centers on two alternative views. According to the specialized mechanism view, there is a specific module in the brain responsive only to faces and not to other objects (e.g., Kanwisher, McDermott, & Chun, 1997). According to the alternative view, the same general mechanism underlies all of object recognition, and what makes face recognition special is our expertise with human faces (e.g., Gauthier & Tarr, 1997). Given that the two paradigms—savings and false recognition—that have provided clear evidence for actor-linked spontaneous inferences rely on faces, this debate on face recognition is certainly informative about research questions concerning the formation and maintenance of person representations. Borrowing from the paradigms used in face recognition research to study the mechanisms underlying the effects described in this article may be fruitful for further specifying person perception mechanisms.

Conclusions

Unintentional trait inferences triggered by single behaviors are bound more strongly to representations of the actors who performed the behaviors than to representations of other people presented at the same time. Differences in the time that attention is allocated to actors at encoding cannot account for this effect. Although these inferences are unintentional, they become integrated into the actor's representation. In fact, the implications of these inferences for the actor representations seem to be the same as the implications of explicit traits; they are indistinguishable after a week. We suggest that simple associative mechanisms cannot account for these effects. The modified false recognition paradigm used in these studies yields robust effects and should be useful in further research on the cognitive processes that produce integrated person representations.

References

Bassili, J. N., & Smith, M. C. (1986). On the spontaneity of trait attributions: Converging evidence for the role of cognitive strategy. *Journal of Personality and Social Psychology, 50,* 239–246.

Brown, R. D., & Bassili, J. N. (2001). Spontaneous trait associations and the case of the superstitious banana. *Journal of Experimental Social Psychology, 38,* 87–92.

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.

Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology, 38,* 889–906.

Frith, C. D., & Frith, U. (1999, November 26). Interacting minds—a biological basis. *Science, 286,* 1692–1695.

Gallagher, H. L., & Frith, C. D. (2003). Functional imaging of "theory of mind." *Trends in Cognitive Sciences, 7,* 77–83.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: Exploring mechanisms for face recognition. *Vision Research, 37,* 1673–1682.

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

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience, 17,* 4302–4311.

Lillard, A., & Skibbe, L. (in press). Theory of mind: Conscious attribution and spontaneous trait inference. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.). *The new unconscious.* New York: Oxford University Press.

Lundqvist, D., Flykt, A., & Öhman, A. (1998). *The Karolinska directed emotional faces.* (Available from the psychology section, Department of Clinical Neuroscience, Karolinska Institute, Stockholm, Sweden)

Mae, L., Carlston, D. E., & Skowronski, J. J. (1999). Spontaneous trait transference to familiar communicators: Is a little knowledge a dangerous thing? *Journal of Personality and Social Psychology, 77,* 233–246.

Malle, B. F. (in press). Folk theory of mind: Conceptual foundations of human social cognition. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious.* New York: Oxford University Press.

Rosenthal, R., & Rubin, D. B. (1982). Comparing effect sizes of independent studies. *Psychological Bulletin, 92,* 500–504.

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.

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 efficiency of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology, 39,* 549–562.

Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352–373.
Uleman, J. S. (1988). [Trait and gist inference norms for over 300 potential trait-implying sentences]. Unpublished raw data.
Uleman, J. S. (1999). Spontaneous versus intentional inferences in impression formation. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 141–160). New York: Guilford Press.
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.
Wellman, H. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.

Received July 31, 2003
Revision received April 26, 2004
Accepted May 12, 2004 ■

Correction to Sagarin et al. (2002)

In the article "Dispelling the Illusion of Invulnerability: The Motivations and Mechanisms of Resistance to Persuasion," by Brad J. Sagarin, Robert B. Cialdini, William E. Rice, and Sherman B. Serna (*Journal of Personality and Social Psychology*, 2002, Vol. 83, No. 3, pp. 526–541), on p. 535, second column, in the third sentence of the *Demonstrated vulnerability treatment* section, all scale labels should have been included. The sentence should read as follows:

> The initial question asked them to indicate how convincing they found the ad on a 7-point scale labeled *not at all convincing* (0), *somewhat convincing* (1), *fairly convincing* (2), *convincing* (3), *quite convincing* (4), *very convincing* (5), and *extremely convincing* (6).