Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Experimental Social Psychology 39 (2003) 549–562

Journal of Experimental Social Psychology

www.elsevier.com/locate/jesp

The efficiency of binding spontaneous trait inferences to actors' faces

Alexander Todorov[a,*] and James S. Uleman[b]

[a] Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010, USA
[b] Department of Psychology, 6 Washington place, New York University, New York, NY 10003, USA

Received 5 June 2002; revised 22 February 2003

Abstract

Three experiments tested the hypothesis that the process of binding spontaneous trait inferences (STIs) to actors' representations is relatively independent of attentional resources. Participants were presented with faces paired with single behaviors. Binding of STIs to actors was revealed by a higher false recognition of implied traits paired with actors' faces than of implied traits randomly paired with other familiar faces. This effect replicated when each face–behavior pair was presented for 2 s (Experiment 1), when the processing of the information was shallow (Experiment 2), and when participants performed a secondary task during the presentation of behaviors (Experiment 3). Experiment 4 showed that explicit on-line trait judgments of the actor, but not explicit behavior judgments, predicted the false recognition of implied traits in the context of the actor. The possibility that the process of binding STIs to actors' representations is automatic is discussed.

Keywords: Person representations; Automaticity; Spontaneous trait inferences; False recognition

Introduction

Imagine yourself out for a midnight stroll in the city, occasionally sighting strangers doing trait-diagnostic things. This one up ahead mutters to himself and greets lampposts politely. That one whistles a sprightly tune to herself. Without even realizing it, you make inferences about each (*mad* and *cheerful*), and stroll on. Contrast this with a midday dash to a lunch appointment when you are preoccupied with what you are going to say when you get there. You pass hundreds of strangers and notice most of them, including these two. One mutters to himself as before, and another whistles a tune, as before. Assuming that you infer *mad* and *cheerful* as before, two questions arise. First, have you made inferences about the people or merely about their behavior? Is the man mad, or is he just behaving like a madman? Is she a cheerful person, or just behaving cheerfully, perhaps only for the moment? Second, will you remember *who* did (or was) what? In your dash to lunch, will you remember which person among all those strangers was cheerful, or merely that someone was cheerful? Keeping straight who did (or was) what seems much more likely on your midnight stroll, when you were not preoccupied and had fewer strangers to remember. Will you keep them straight at midday?

We can restate these two questions more precisely. First, are spontaneous trait inferences (STIs) from people's behavior merely about the behavior, or are they also about the actor performing the behavior? Second, how efficiently are such trait inferences linked or bound to particular appropriate actors? The first three studies in this paper address this second question, using a new false recognition paradigm introduced by Todorov and Uleman (2002). The last study addresses the first question.

Many experiments have shown that simply reading about or observing a single behavior by another person is sufficient to trigger a trait inference, even when the observer does not intend to make an inference or to form an impression of the actor (for a review see Uleman, Newman, & Moskowitz, 1996). That is, people make spontaneous trait inferences. Based on this evidence, it has been commonly assumed that spontaneous behavior-to-trait inferences are an automatic process (Bargh, 1994), but that these inferences are limited to the meaning of the behavior and do not have further implications for the

* Corresponding author. Fax: 1-609-258-1113.
E-mail address: atodorov@princeton.edu (A. Todorov).

0022-1031/$ - see front matter
doi:10.1016/S0022-1031(03)00059-3

actor (Bargh, 1994; Bassili, 1989; Park, 1989; Wyer & Srull, 1989). In other words, such unintended inferences result in higher accessibility of trait concepts in memory but are not linked to the representation of the actor.

However, two recent experimental paradigms have provided convergent evidence strongly suggesting that STIs are about the actor. Both the savings and the false recognition paradigms use faces as critical stimuli. In the savings paradigm (Carlston & Skowronski, 1994; Carlston, Skowronski, & Sparks, 1995), participants are presented with faces paired with behavioral paragraphs implying traits. Later they are asked to learn face–trait pairs. Participants are faster at learning faces paired with traits that were implied in the paragraph presented with the face, than faces paired with new traits, or paired with traits implied in a paragraph about a different face. This facilitation of learning suggests that the trait was previously inferred and linked to the actor's face.

In the false recognition paradigm (Todorov & Uleman, 2002), participants are presented with faces paired with trait implying behavioral sentences (e.g., "Mary solved the mystery half-way through the book"). Later in a recognition test, participants are presented with the faces paired with trait words (e.g., "clever") and asked to decide whether they saw the word in the sentence presented with the face. Participants are more likely to falsely recognize implied traits when these traits are paired with actors' faces than when they are paired with other familiar faces presented earlier in the experiment.[1]

The evidence from the false recognition and savings paradigms suggests that STIs are linked to representations of the actors. However, none of the previous studies using these two paradigms has addressed the automaticity of these inferences. For example, it is possible that behavior-to-trait inferences are automatic (e.g., Bargh, 1994; Lupfer, Clark, & Hutcherson, 1990; Winter, Uleman, & Cunniff, 1985), but that the process of binding these inferences to actor representations depends on the availability of cognitive resources and could be easily blocked by manipulations that interfere with the processing of the behavioral information. The objective of the present research was to address the automaticity of binding STIs to actors.

The robustness of actor-linked STIs

Although the previous studies in the savings and false recognition paradigms have not addressed the automaticity of actor-linked STIs directly, several experiments have shown that these inferences are fairly robust. For example, the savings effect (faster learning of face and implied-trait pairs) is unaffected by encoding instructions (Carlston & Skowronski, 1994). This effect seems to be the same for familiarization, memory, and impression formation instructions. In the false recognition paradigm, Todorov and Uleman (2002, Experiment 4) have arrived at similar conclusions, although they showed that memory and impression instructions trigger different processes. Moreover, even under conditions of information overload when participants were presented with 120 unfamiliar faces, each paired with a different behavior, participants linked STIs to the actor representation (Todorov & Uleman, 2002, Experiments 5 & 6).

In both paradigms, actor-linked STIs are not dependent on behavior retrieval. In the false recognition paradigm, the false recognition of implied traits occurs even if participants do not recall or do not recognize the behavior (Todorov & Uleman, 2002, Experiments 5 & 6). Similarly, the savings effect occurs even if participants do not recognize the behavioral paragraphs and the learning task is a week later than the presentation of faces and paragraphs (Carlston & Skowronski, 1994; Carlston et al., 1995).

However, in all previous studies the presentation time of faces paired with behavioral information has been sufficiently long (ranging from 5 to 10 s) to allow participants to rehearse the information or develop some consistent strategy for memorizing the information. Also, during this initial presentation participants' only task was to memorize or familiarize themselves with the information. None of the previous experiments used a simultaneous secondary task to tax participants' cognitive resources, or a low-level processing goal that produces shallow processing. Thus, it is not clear whether actor-linked inferences occur when cognitive resources are severely constrained.

How efficiently are STIs bound to actors?

The first three experiments used the false recognition paradigm (Todorov & Uleman, 2002) to examine the efficiency of STI-actor binding. Efficiency is one characteristic of automatic processes (Bargh, 1994) and includes relative independence of attentional resources. Manipulations that deplete attentional resources include rapid presentation of stimuli (e.g., Bargh & Thein, 1985) and a concurrent cognitive load (e.g., Gilbert & Osborne, 1989). To the extent that a process is robust with respect to such manipulations, the process is efficient. Experiment 1 reported below used rapid pace presentation of face–behavior stimuli, and Experiment 3 put participants under a cognitive load, rehearsing 6-digit numbers. A process can also be efficient in the sense that it occurs even under shallow processing goals

[1] To make the recognition task believable, some sentences contain trait words (e.g., "Mary was clever and solved..."). The rate of correct recognition of presented traits is generally higher than the rate of false recognition of implied traits, indicating that participants are paying attention to the task. If the rates of correct recognition and false recognition were identical, that would suggest that participants perform at a chance level.

(Craik & Lockhart, 1972). In Experiment 2, participants were told that the experiment was about how people process grammatical information in the context of faces and were asked to count the nouns in each sentence.

All three manipulations—rapid pace, cognitive load, and shallow processing—should lead to less accurate recognition performance. Compared with control conditions, participants should be both more likely to falsely recognize traits (whether or not these traits are paired with the actor's face) and less likely to correctly recognize presented traits. More importantly, if binding of STIs to actor representations is relatively independent of attentional resources, participants should be more likely to falsely recognize implied traits paired with actors' faces than implied traits randomly paired with familiar faces, even under these conditions.

The nature of actor-linked spontaneous inferences

The evidence from the false recognition and savings paradigms clearly shows that people engage in spontaneous inferences and that these inferences are linked to the actor representations. However, the nature of these inferences is ambiguous. We (Todorov & Uleman, 2002) have assumed that these are trait inferences about the actor. However, a simpler assumption is that these are inferences about the meaning of the behavior that are merely associated with actor representations. Indeed, the false recognition and the savings findings do not discriminate between these two assumptions.

Experiment 4 was designed to test whether the false recognition of implied traits reflects spontaneous trait inferences *about* the actor rather than merely behavioral categorizations associated with the actor representations. We measured on-line trait judgments of the actor ("Is Robert a honest person?") and on-line behavior judgments ("Is this a honest behavior?"). If spontaneous inferences are inferences about the actor, explicit on-line judgments of the actor, but not on-line behavior judgments, should predict the probability of false recognition of implied traits.

Experiment 1

In all previous experiments using the false recognition paradigm, each face–behavior pair was presented either for 10 s or for 5 s. In this experiment, the presentation time was reduced to 2 s. This, coupled with the relatively large number of face–behavior pairs (60), reduces the probability of rehearsing the behaviors and other deliberate strategies to memorize the information. The experiment also included a condition in which participants self-paced the presentation of face–behavior pairs, allowing us to estimate the average time that participants prefer to spend studying each face–behavior pair.

Relative to the self-pace condition, participants in the fast-pace presentation condition should be both more likely to falsely recognize traits and less likely to correctly recognize presented traits. More importantly, if actor-linked STIs are efficient, two seconds should be enough for participants to make a trait inference and to associate this inference with the actor. Specifically, they should be more likely to falsely recognize implied traits paired with the actors' faces than implied traits randomly paired with other familiar faces.

In all experiments, the false recognition control trials consisted of familiar faces randomly paired with implied traits. Thus, differences between false recognition of implied and control traits cannot be accounted for in terms of familiarity of either faces or traits. In fact, Todorov and Uleman (2002) found that such control traits are more likely to be falsely recognized than novel control traits. This suggests that randomly paired implied traits provide a more stringent baseline than novel control traits do, against which to compare the false recognition of implied traits paired with actors' faces.

Method

Participants

Twenty-seven undergraduates from the Department of Psychology at New York University participated in the study for partial course credit. Participants were randomly assigned to the two experimental conditions.

Stimulus material

Sixty 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. The sentences represented a wide range of behaviors, implying different traits. In the recognition test, each trait word was presented only once. The black and white pictures were of college students from universities other than New York University. To avoid any confounding of faces (e.g., positive expressions) with particular types of behaviors (e.g., positive behaviors), the sentences were randomly paired with faces. All face–behavior pairs were used in Todorov and Uleman (2002, Experiments 5 & 6).

Procedures

Participants were told that this was a study of how people remember information. Participants worked individually in soundproof cubicles and instructions were presented by a computer. Participants were told that the experiment consisted of two parts. In the first part, they would be shown pictures of people with information about them and in the second part their memory would be tested. The type of memory test was not specified in the instruction. The experiment started with a practice trial and, if everything was clear to the participant, continued with the study phase. Each participant was presented with 60 stimulus pairs (trials) of a face and a behavioral sentence. The order of the 60 trials was randomized for each participant by the computer. In the fast-pace presentation condition, each trial (a face plus a sentence) was presented for 2 s. In the self-pace presentation condition, participants decided when to end the trial presentation. The time delay between trials was 2 s. In 20 of the trials, the sentences contained the trait implied by the behavior. In the remaining 40 trials, the sentences only described behaviors.

After the 60 study trials, participants were told that in the second part, they would be presented with the faces from the first part of the experiment, each accompanied by a single word. Their task was to decide whether they saw the word in the sentence about the person. The participant's task was to press the "OLD" key on the keyboard (the "M" key, labeled "OLD") if they believed that they saw the word in the study phase, or the "NEW" key (the "X" key, labeled "NEW") if they believed that they did not see the word. Participants were asked to work as quickly as possible. In order to familiarize them with the task, they had two practice trials. The face and the sentence from the practice trial in the study phase were presented for 5 s. Then the face was presented with a trait word, which was part of the sentence. Before this trial, participants were reminded about the "OLD"/"NEW" decision. All participants pressed the "OLD" key. After the trial, participants were told that this was the correct decision and then were presented with a face and a trait 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.

In the test phase, participants were presented with the 60 faces from the study phase each paired with a trait word. The order of the test trials was randomized for each participant by the computer. The trait word was presented below the face. Each test trial stayed on the screen until the participant's response. The next trial followed after 2 s. In 20 trials, participants were presented with the faces from the filler sentences, which contained trait words, paired with the trait words actually presented. The remaining 40 faces were randomly divided into two groups of 20 faces. In one of the groups, each face was randomly paired with a trait implied in a sentence about another face from the same group. These 20 face–trait pairs served as control trials. In the other group, each face was paired with the implied trait. These groups of face–trait pairs were counterbalanced in Experiments 5 and 6 reported by Todorov and Uleman (2002). Since there was no effect of this counterbalancing factor, and that was true in the other four experiments reported in that paper, the counterbalancing factor was dropped from the current experiment. The overall design was a mixed 3 (Trait: presented vs. implied vs. control) × 2 (Pacing: fast-pace vs. self-pace) ANOVA with the first factor within-subjects and the second between-subjects.

Analyses

All analyses were conducted at both the level of participants and the level of stimuli. The analysis at the level of stimuli treats participants as variables and stimuli (face–behavior pairs) as cases. To the extent that the observed effects are robust and generalize across participants and stimuli, both types of analyses should agree.

Given the specific hypotheses, two main analyses were performed. The first analysis was conducted as a manipulation check of the procedures. It collapsed across the false recognition of implied and control traits, and compared false recognition and correct recognition in the two experimental conditions. This 2 (Trait Presentation: presented vs. implied) × 2 (Pacing) ANOVA was performed to establish that overall accuracy decreased under the fast-paced presentation. In all experiments, correct recognition was significantly higher than the false recognition. Because these effects are not theoretically relevant, we do not report these main effects. We do, however, report the interaction of Trait Presentation and the other experimental condition.

A second 2 (Implied Trait: actor vs. control) × 2 (Pacing) ANOVA compared rates of false recognition of implied and control traits as a function of the experimental condition. This is the most critical analysis, because it tests the hypotheses that implied traits are associated with the actor and that this effect is not dependent on presentation pace.

Results

Manipulation checks

In the self-pace condition, the average time for reading a sentence paired with a face was 6.29 s.[2] This was significantly longer than 2 s, $t(12) = 8.93$, $p < .001$.

[2] Todorov and Uleman (2002, Experiment 4) showed that participants under memory instructions were less likely to falsely recognize implied traits when each face–behavior pair was presented for 10 s than for 5 s. Consistent with this finding, correlational analysis at the level of the stimuli in the self-pace condition showed that the probability of false recognition of implied traits was highly correlated with reading times, $r(20) = -.61$, $p < .005$. The longer the reading time, the less likely the false recognition of the implied trait was. These findings suggest that participants made fast trait inferences, but that the additional presentation time allowed them to create a more detailed behavioral representation. Correspondingly, they were less likely to confuse the inference with the perception of the trait. Importantly, the manipulation of the presentation time in Experiment 4 of Todorov and Uleman (2002) did not affect the false recognition of control traits. Similarly, the correlation in the self-pace condition between the false recognition of control traits and reading times was not significant ($r = .24$).

The analysis collapsing across the false recognition of implied and control traits showed a reliable interaction of Trait Presentation and Pacing, $F(1,25) = 13.76$, $p < .001$, indicating that the fast-pace presentation led to both a lower rate of correct recognition of presented traits and a higher rate of false recognition of implied traits. The interaction was also significant at the level of the stimuli, $F(1,58) = 41.64$, $p < .001$. Participants in the 2-s condition were less likely to correctly recognize presented traits ($M = .66$, $SD = .12$) than participants in the self-pace condition ($M = .77$, $SD = .12$), $t(25) = 2.35$, $p < .027$, and $t(19) = 3.41$, $p < .003$ for stimuli. Participants in the 2-s condition were also more likely to falsely recognize traits ($M = .43$, $SD = .13$) than participants in the self-pace condition ($M = .28$, $SD = .16$), $t(25) = 2.72$, $p < .012$, and $t(39) = 6.40$, $p < .001$ for stimuli.

False recognition of traits

As shown in Fig. 1, the analysis of the false recognition responses showed that across both conditions participants were more likely to falsely recognize actor implied traits ($M = .46$, $SD = .20$) than control traits ($M = .24$, $SD = .16$), $F(1,25) = 42.48$, $p < .001$. This effect was also significant at the level of the stimuli, $F(1,38) = 24.25$, $p < .001$. The Implied Trait × Pacing interaction was not significant for either analysis, Fs < 1. The lack of significant interaction indicates that the effect was equivalent for both pacing groups. If anything, in the 2-s condition the effect size ($r = .87$) was slightly larger than the effect size ($r = .74$) in the self-pace condition.

Discussion

As expected, the fast-pace presentation of face–behavior pairs reduced the overall accuracy of participants. Participants in this condition were both less likely to recognize presented traits and more likely to falsely recognize other traits. More importantly, the reduced presentation time did not affect the difference between the false recognition of implied traits with actors and the false recognition of control traits. Although this presentation rate was three times faster than the self-pace rate, participants in the fast-pace condition were just as likely as participants in the self-pace condition to associate the implied trait with the actor. This suggests that STIs and actors were linked on-line within 2 s. Note that these 2 s include the comprehension of the behavioral information and the integration of this information with the perception of the face.

Experiment 2

Experiment 2 manipulated the encoding goals of participants. They were either given memory instructions as in Experiment 1 or were asked to count the nouns in each sentence, ostensibly as part of a task on grammatical processing in the context of faces. Because the latter instruction does not require elaborating on the behavioral meaning or processing the actor's face, it should lead to shallow information processing. In explicit memory tasks, where participants are asked to

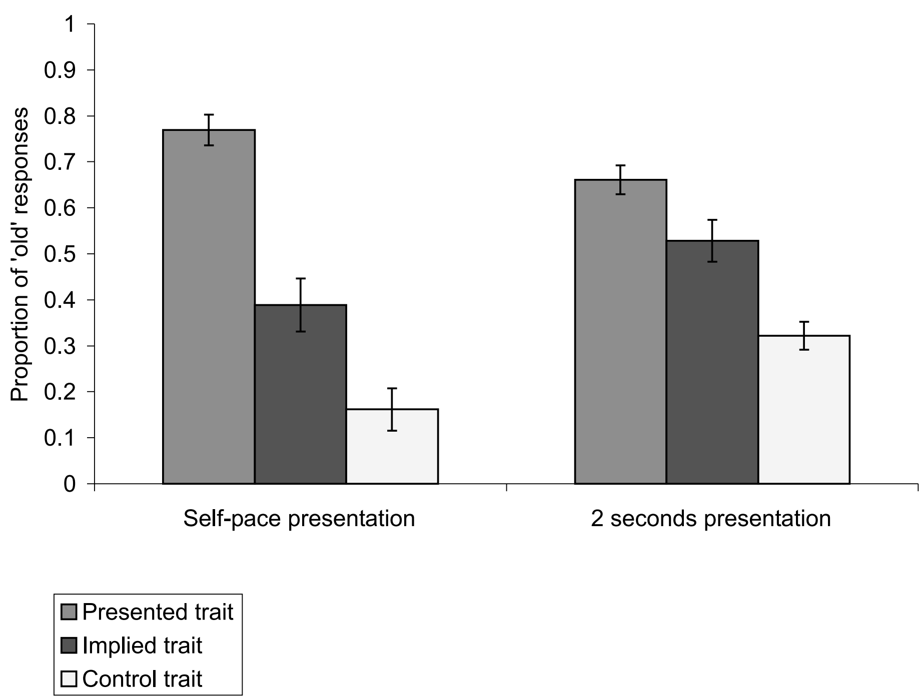

Fig. 1. Mean proportion of 'OLD' recognition responses as a function of face–trait pairing, type of trait, and experimental condition (Experiment 1).

recall or recognize information, shallow processing results in poor memory for the material (e.g., Craik & Lockhart, 1972). One explanation for this effect is that shallow processing does not facilitate the formation of links between the studied materials that can serve as retrieval cues in a later test (e.g., Craik & Tulving, 1975).

This is especially relevant in the false recognition paradigm which tests for links between actors and inferences. If such links can be formed only after careful elaborative processing of the information, then shallow processing should prevent the formation of such links. In fact, there is some evidence in the literature that STIs do not occur under conditions of shallow processing. Using the cued recall paradigm, in which implied traits serve as retrieval cues for studied behavioral sentences, Uleman and Moskowitz (1994) showed that STIs virtually disappear when participants' task does not involve processing the meaning of the sentences (e.g., checking each word for a particular phoneme).

The shallow processing instructions also control the type of learning or rehearsal strategies that participants might use. Under memory instructions participants can pursue a variety of strategies for memorizing the material, and these strategies might include elaborating on the actor's personality. In addition, it is possible that faces paired with behavioral information trigger automatic impression formation goals (see Bargh, 1990, and Bargh & Barndollar, 1996 on automatic goals). If this is the case, then the trait inferences that occur may not be spontaneous after all, despite the memory instructions. Todorov and Uleman (2002, Experiment 4) provided evidence against this hypothesis by showing that memory and impression instructions are affected differentially by the presentation time of face–behavior stimuli. Whereas doubling this time did not affect the false recognition of implied traits under impression instructions, it substantially reduced the false recognition under memory instructions. However, it is possible that some participants in the memory condition pursued impression goals. Thus a second objective of this experiment was to provide additional, more direct evidence against the hypothesis that associating inferred traits with actors is driven by automatic impression formation goals.

If participants associate the inferred trait with the actor, they should be more likely to falsely recognize implied than control traits in both memory and shallow processing conditions. At the same time, shallow processing should lead to lower recognition accuracy.

Method

Participants

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

Procedures

In the shallow processing condition, participants were told that this was a study of how people process grammatical information in the context of faces. They were told that on each trial they would see a face and a sentence, and that their task would be to count the nouns in each sentence. Participants were also told to read the information carefully because they would be asked additional questions at the end of the experiment. After a practice trial and a correct response to it, the study phase of the experiment began. Each face–behavior pair was presented for 5 s and immediately followed by a noun judgment. Participants pressed the number on the keyboard corresponding to the number of nouns in the sentence. The order of the trials was randomized for each participant, and the delay between trials was 2 s. In the memory condition, participants followed the same procedures as in Experiment 1 except that each face–behavior pair was presented for 5 s.

After 60 face–behavior pairs, participants took the word recognition test using the same procedures as in Experiment 1. The overall design was a mixed 3 (Trait: presented vs. implied vs. control) $\times$ 2 (Goal: counting nouns vs. memory) ANOVA with the first factor within-subjects and the second between-subjects. The same types of analyses were performed as in Experiment 1.

Results

Manipulation checks

The analysis collapsing across the false recognition of actor-implied and control-implied traits revealed a significant Trait Presentation (presented vs. implied) $\times$ Goal interaction, $F(1,36) = 10.06$, $p < .003$, also significant at the level of the stimuli, $F(1,58) = 12.57$, $p < .001$. Participants with the memory goal were more likely to correctly recognize presented traits ($M = .71$, $SD = .17$) than those counting nouns ($M = .59$, $SD = .15$), $t(36) = 2.15$, $p < .039$. This was also significant at the level of stimuli, $t(19) = 2.43$, $p < .025$. Participants with a memory goal were also less likely to falsely recognize traits ($M = .42$, $SD = .18$ vs. $M = .34$, $SD = .16$ respectively), $t(36) = 1.43$, $p = .16$. Although the latter effect was not significant, it was in the predicted direction and significant at the level of the stimuli, $t(39) = 2.59$, $p < .014$.

False recognition of traits

Participants were more likely to falsely recognize actor-implied traits ($M = .48$, $SD = .22$) than control traits ($M = .27$, $SD = .15$), $F(1,36) = 65.46$, $p < .001$, for participants and $F(1,38) = 28.12$, $p < .001$, for stimuli. However, this effect was qualified by a significant Implied Trait (actor vs. control) $\times$ Goal interaction, $F(1,36) = 4.99$, $p < .032$, for participants, and $F(1,38) = 4.19$, $p < .048$, for stimuli (Fig. 2, left and

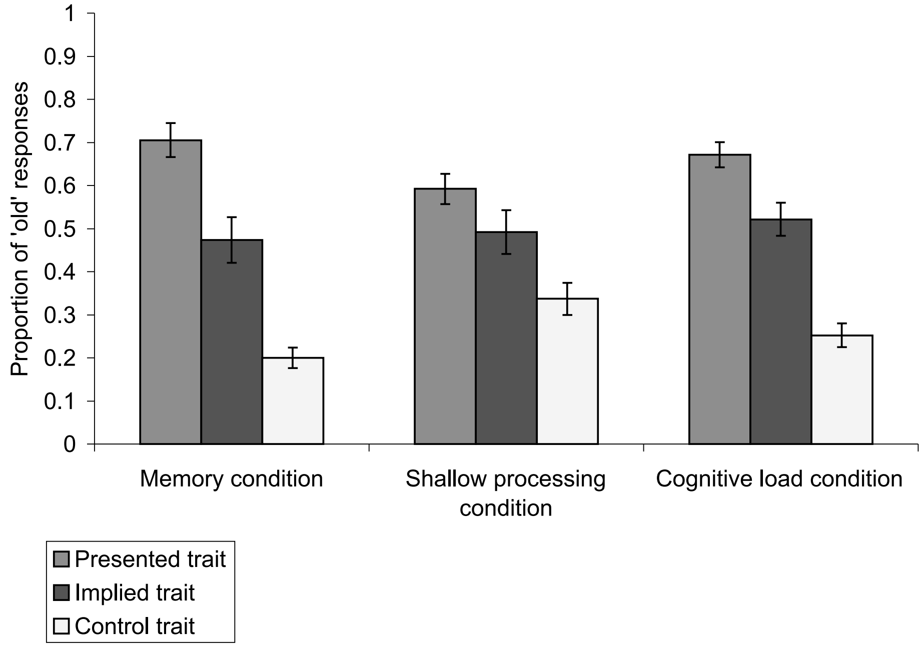

Fig. 2. Mean proportion of 'OLD' recognition responses as a function of face–trait pairing, type of trait, and experimental condition (Experiment 2, left and center; Experiment 3 on the right).

center). The effect of higher false recognition of implied traits relative to control traits was weaker in the shallow processing condition than in the memory condition. Nevertheless, although the effect was reduced in this condition, it remained reliable, $t(18) = 4.43$, $p < .001$, for participants, and $t(38) = 3.03$, $p < .004$, for stimuli. The effect size in this condition ($r = .72$) was comparable to the effect size in the memory condition ($r = .85$).

Accuracy of noun counting

In the shallow processing condition, participants' decisions about the number of nouns were correct on 60.3% of the trials, indicating that the task was difficult. It is possible that the false recognition of actor-implied traits was higher than that of control traits only when participants were inattentive to the primary task (counting the nouns in the sentence). That is, when participants did not pay attention to the primary task, they would have been more free to pursue other processing goals. To rule out this possibility that our effect depends on inattention to noun counting, the data in the shallow processing condition were analyzed as a function of accuracy in counting nouns. If participants associate the inferred trait with the actor efficiently, the accuracy of noun counts should not affect the difference between the false recognition of actor-implied and control traits. In contrast, if this difference depends on inattention to the primary task on some trials, the false recognition of actor-implied traits should be higher than the false recognition of control traits only on trials when participants got the noun count wrong.

The data were analyzed in a 2 (Noun Count: correct vs. incorrect) $\times$ 2 (Implied Trait: actor vs. control) repeated measures ANOVA. The only significant effect was the effect of Implied Trait, $F(1, 18) = 17.89$, $p < .001$; all other Fs < 1. Thus, the false recognition effect under shallow processing reported above did not occur because of occasional inattention to the noun counting task.

Discussion

The shallow processing of behavioral information had similar effects on recognition performance as the fast-pace presentation in Experiment 1. Participants in the shallow processing condition were more likely to falsely recognize-implied traits, and less likely to correctly recognize presented traits than participants in the memory condition. More important, in both conditions participants were more likely to falsely recognize actor implied than control traits, although this effect was somewhat reduced in the shallow processing condition. The difference between the false recognition of actor-implied traits and control-implied traits in the shallow processing condition cannot be explained in terms of inattention to the primary task of counting nouns.

Thus, even when the participants' goal did not involve processing the behavioral information in relation to the actor's face, participants associated the implied and inferred trait with the actor. This suggests that the occurrence of actor-linked STIs requires minimal processing of behavioral information. Interestingly, this finding is inconsistent with findings from the cued recall paradigm (Uleman & Moskowitz, 1994) that failed to provide conclusive evidence that STIs are linked to the actor's representation. The two major differences between the cued recall and the false recognition paradigms are that the latter uses faces as critical stimuli and exploits both actor-to-trait and trait-to-actor links. Cued recall does not use such strong individuating information as faces and only exploits trait-to-actor links.

This experiment also provided evidence against the hypothesis that some of the actor-linked STIs are driven by automatic impression formation goals. Participants' goals in the shallow processing condition clearly did not entail making any inferences about the actor or elaborating on the information in any way. In fact, this was confirmed by the poor memory performance in this condition. Nevertheless, participants formed links between the actor's face and the implied trait.

Experiment 3

The first two experiments showed that actor-linked STIs occurred when the presentation pace of face-behavior stimuli was rapid and when the behavioral information was processed shallowly. Experiment 3 used an additional manipulation intended to interfere with strategic or controlled but not automatic processing. Participants processed the face–behavior stimuli under cognitive load. Each trial consisted of a presentation of a 6-digit number followed by the presentation of a face–behavior pair and then a number recognition task. Participants had to rehearse the 6-digit number while reading the behavioral information, and then recall it. Cognitive load is one of the most frequently used manipulations for testing the efficiency of a cognitive process (Bargh, 1994). Efficient processes should not be affected by the depletion of attentional resources under a cognitive load such as digit rehearsal.

Method

Participants

Twenty-one undergraduates from the Department of Psychology at NYU participated in the study for partial course credit.

Procedures

Participants were told that this was a study of how people process two types of information simultaneously. They were told that each trial would start with a 6-digit number presented for 2 s, followed by a face and a sentence presented for 5 s, followed by a number recognition task. Their task was to start rehearsing the number and, while rehearsing the number, to read the information presented with the face on each trial. Each face–behavior pair was followed by a number recognition task. The delay between the face–behavior pair and the recognition task was 500 ms. On half of the 60 trials the number was the same as the one presented before the face–behavior pair, and on the other half the number was different. The participants' task was to press the "OLD" key (the "M" key on the keyboard) if the number was the same or the "NEW" key (the "X" key on the keyboard) if the number was different. Participants were told to read the information carefully because their memory would be tested later in the experiment. The memory test was not specified further.

Each of the 60 study trials consisted of a 6-digit number followed by a face–behavior pair followed by a number recognition task. Participants started with a practice trial and were given feedback for this trial. Then they continued with two new practice trials, and after the task was clear, they proceeded with the study phase of the experiment. The order of the trials was randomized for each participant, and the delay between trials was 2 s. The 60 6-digit numbers were randomly generated from a uniform distribution (100000, 999999). For example, one of the numbers was 783891. For the number recognition test, half of the numbers were changed to a different number by replacing one of the digits with a different digit close to the original digit (e.g., 783891 changed to 784891). After the 60 trials, participants were given the word recognition test. The procedures for this test were the same as in Experiments 1 and 2.

Experiment 3 did not include a between-subjects control condition because participants were drawn from the same population, during the same semester, as in Experiments 1 and 2.

Results

Manipulation checks

Because this experiment did not include a control condition, we compared the participants' performance with the self-pace condition in Experiment 1 and the memory condition in Experiment 2. Relative to both conditions, participants under cognitive load were less likely to correctly recognize presented traits ($M = .67$, $SD = .13$) and more likely to falsely recognize implied traits ($M = .39$, $SD = .13$); $F(1, 32) = 14.64$, $p < .001$, for the interaction with the self-pace condition at the level of participants, and $F(1, 58) = 41.87$, $p < .001$, at the level of stimuli; and $F(1, 38) = 3.82$, $p < .058$, for the interaction with the memory condition at the level of participants, and $F(1, 58) = 5.04$, $p < .029$, at the level of stimuli.

False recognition of traits

Participants under load were more likely to falsely recognize actor implied traits ($M = .52$, $SD = .18$) than control-implied traits ($M = .25$, $SD = .13$), $t(20) = 7.78$, $p < .001$, effect size $r = .82$ (see Fig. 2, right side), as predicted. This effect was also significant at the level of the stimuli, $t(38) = 7.93$, $p < .001$.

Accuracy of number recognition

The average accuracy on the number recognition task was .82. This proportion was significantly higher than the chance level of .50, $t(20) = 13.55$, $p < .001$, but significantly lower than the perfect performance of 1, $t(20) = 7.37$, $p < .001$. As in Experiment 2, the data were also analyzed as a function of the response to the primary (number recognition) task. The only significant effect found by this 2 (Response to number recognition task: correct vs. incorrect) $\times$ 2 (Implied Trait: actor vs. control) repeated measures ANOVA was the effect of trait, $F(1, 20) = 21.73$, $p < .001$; other Fs < 1. Thus, the predicted effect did not depend on inattention to the cognitive load task.

Discussion

Consistent with the findings of Experiments 1 and 2, participants were more likely to falsely recognize actor-implied than control-implied traits. And as in Experiment 2, this difference cannot be explained in terms of inattention to the other task (counting nouns in Experiment 2 and number recognition in the current experiment). Even when participants had to perform two tasks simultaneously, they were able to link the inferred trait to the actor's representation. This finding, coupled with the same finding in the fast-pace presentation condition in Experiment 1, suggests that making actor-linked STIs is a highly efficient process.

Experiment 4

Are STIs about the actor, or merely about the behavior and linked to the actor? Are people inferring that Mary is a clever person, or merely did something clever? If STIs are about the actor rather than merely about the behaviors, then measures of person inferences, but not of behavior inferences, should predict the probability of false recognition of implied traits. Todorov and Uleman (2002) provided evidence that explicit person judgments on the implied traits predicted the false recognition of these traits in the context of the actor's face. However, this study did not include a comparable measure of behavior inferences.[3] Given that person inferences are grounded in the meaning or trait diagnosticity of the behavior, it is plausible that person inferences are confounded with behavior inferences. The objective of Experiment 4 was to partial out the effects of person inferences and behavior inferences on false recognition.

In Todorov and Uleman (2002), trait judgments of the actor predicted the false recognition of implied traits even though the procedures used to collect the judgments and to measure the false recognition were clearly different. The actor judgments were measured on 11-point scales in a questionnaire study, in which each participant made three different trait judgments of the actor (on the implied trait, on a novel unrelated trait, and on an antonym of the implied trait). In the current study, we measured both on-line judgments of the actor and on-line judgments of the behavior with procedures similar to the experimental procedures of Experiments 1, 2, and 3. Participants were asked to make either person judgments or behavior judgments. In the case of person judgments, each face–behavior pair was presented for 5 s, and immediately after this presentation participants were asked to make a Yes/No trait judgment of the actor. In the case of behavior judgments, each behavior was presented for 5 s, and participants were asked to make a Yes/No trait judgment of the behavior (whether the trait characterized the behavior).

If the false recognition of implied traits reflects spontaneous person inferences made during the presentation of the faces with behaviors, then it should be predicted by explicit person judgments. Alternatively, if the false recognition reflects spontaneous behavior inferences, it should be predicted by explicit behavior judgments. The critical test of these hypotheses requires mediation or partial correlation analyses because prominent models of dispositional inferences (Gilbert & Malone, 1995; Trope, 1986) assume that dispositional (person) inferences follow initial behavior inferences. Further, these inferences are highly correlated (e.g., Trope & Alfieri, 1997). Thus, person inferences can be correlated with false recognition by virtue of their correlation with behavior inferences. If this is the case, the correlation between person judgments and false recognition should be reduced substantially after the effect of behavior inferences is controlled for. In contrast, if false recognition reflects spontaneous person inferences, this correlation should not be affected. Moreover, the correlation of behavior inferences and false recognition should drop to zero after the effect of person inferences is controlled for.

In this experiment, we also measured the response times for the person and behavior judgments, and collected additional measures that could be related to the false recognition of implied traits. These measures included the language frequency of the trait words (Francis & Kucera, 1982) and their valence. We expected that none of these variables would be a significant predictor of the false recognition of implied traits,

[3] Todorov and Uleman (2002) included a measure of the probability of generating the trait given the behavior. This probability was calculated during the process of compiling the behavioral sentences (Uleman, 1988). This measure did not correlate with any of the relevant measures either in Todorov and Uleman's studies (2002) or in the current studies, most likely because of the nature of the task (trait generation vs. trait judgment). For example, the correlation of this measure with the behavior judgments measured in the current experiment was. 07.

once the analysis controlled for the trait judgments of the actor.

The analyses could be performed only at the level of stimuli and not at the level of participants, because false recognition measures unintentional inferences and judgments measure intentional inferences about the same actors. An advantage of the analysis at the level of stimuli is that the judgment measures—e.g., the probability of attributing the trait to the actor given the actor's behavior—could be used as predictors of the false recognition measures in each experimental condition of the experiments above and across experimental conditions. If the false recognition of implied traits reflects automatic on-line inferences about the actor, the strength of the relation between the probability of false recognition and the probability of attributing the trait to the actor should not vary as a function of the processing conditions in the experiments above.

Method

Participants

Ninety-eight undergraduates from the Department of Psychology at NYU participated in the study for partial course credit (53 made person judgments and 45 made behavior judgments).

Procedures

In the person-judgment part of the study, participants were told that this was a study about judgments of other people, that on each trial they would be presented with a person's face and a sentence describing the person for 5 s, and that they would make a judgment about the person. Participants were presented with the 40 face–behavior pairs that were used in Experiments 1–3 and did not contain trait words. Each face–behavior pair was followed by a trait judgment. In the false recognition experiments above, each face was paired with the implied trait in the recognition test for 20 of the face–behavior pairs. In the present study, each of these 20 face–behavior pairs was followed by a judgment about the implied trait. For example, a face was presented with the sentence "Richard dusted and vacuumed his room every day." for 5 s, followed by the question "Is Richard a neat person?" If participants believed that the answer was "Yes," they pressed the "YES" key (the "M" key on the keyboard). If they believed that the answer was "No," they pressed the "NO" key (the "X" key on the keyboard). Each of the remaining 20 face–behavior pairs was followed by a judgment about an unrelated trait that was implied about another face. These were the control trials in Experiments 1–3. The order of trials (face–behavior pair followed by Yes/No trait judgment) was randomized for each participant, and the delay between trials was 2 s. Participants started with four practice trials to become familiar with the task.

In the behavior-judgment study, participants were told that this was a study about judgments of behaviors. The procedures were the same as in the person-judgment study with the exception that on each trial, participants were presented with only a behavior and made a Yes/No judgment about the trait meaning of the behavior. For example, the sentence "Richard dusted and vacuumed his room every day." was followed by the question "Is this a neat behavior?"

Analyses

The probability of attributing the trait to the actor was computed for each face–behavior pair.[4] Similarly, the probability of characterizing the behavior as representing the trait was computed for each behavior. These two variables were used as predictors of the false recognition of traits at the level of stimuli. The probability of false recognition of traits was computed for each face–trait pair in each experimental condition of Experiments 1–3 (2 s presentation, 5 s presentation, shallow processing, self-pace, and cognitive load). The analyses also included the probability of false recognition of traits from the relevant conditions of Experiments 5 and 6 reported in Todorov and Uleman (2002). These two experiments used the same procedures as the memory condition in Experiment 2 except that participants were presented with 120 rather than 60 face–behavior pairs and made 120 word recognition decisions.

Results

Manipulation checks

Participants who made person judgments were more likely to attribute the implied trait to the actor (.85) than the control trait (.29), $t(52) = 24.79$, $p < .001$. This difference was also significant at the level of stimuli, $t(38) = 8.36$, $p < .001$. Participants were also faster to make person judgments about the implied trait (1884 ms) than about the control trait (2499 ms), $t(52) = 6.40$, $p < .001$, and $t(38) = 4.89$, $p < .001$, for stimuli. The pattern of data was the same for participants who made behavior judgments. Behaviors were more likely to be judged as exemplifying the implied

[4] The mean response times for positive person and behavior judgments ("Yes" responses) were also calculated for each face–behavior pair. Judgment speed is an indication of the strength of the trait implication. Thus, faster judgments should result in a higher rate of false recognition of implied traits. The correlations between mean judgment times and false recognition rates were not as strong as the correlations of false recognition rates with the judgments themselves. But in all experimental conditions these correlations were negative, indicating that faster judgments were associated with more false recognition errors. Across experimental conditions, the correlation of mean false recognition with the response times for person judgments was −.48, $p < .032$. With the response times for behavior judgments, it was −.44, $p < .055$.

Table 1
Probability of false recognition of implied traits at the level of stimuli and correlations of false recognition with explicit person and behavior judgments as a function of experimental condition

Experimental condition	False recognition	Correlations of false recognition with	
		Person judgments	Behavior judgments
Memory (self-pace)	.39 (.15)	.70[a]	.49[a]
Memory (5 s)	.47 (.17)	.62[a]	.48[a]
Memory (5 s)[b]	.51 (.23)	.63[a]	.50.[a]
Memory (2 s)	.53 (.20)	.70[a]	.60[a]
Memory (Load, 5 s)	.52 (.13)	.58[a]	.38
Shallow processing	.49 (.18)	.60[a]	.46[a]

[a] $p < .05$.
[b] These are data from the relevant conditions of Experiments 5 and 6 reported in Todorov and Uleman (2002). The procedures for these conditions were the same as the procedures in the memory condition (5 s) in Experiment 2 reported here except that participants were presented with 120 rather than with 60 face–behavior pairs.

trait (.86) than the control trait (.26), $t(44) = 12.41$, $p < .001$, and $t(38) = 11.94$, $p < .001$, for stimuli. Similarly, participants were faster to make behavior judgments about the implied (1731 ms) than about the control trait (2115 ms), $t(44) = 8.42$, $p < .001$, and $t(38) = 4.46$, $p < .001$, for stimuli.

Predicting false recognition

The critical question is which judgments predict the false recognition of implied traits.[5] As shown in Table 1, both the probability of attributing the trait to the actor and the probability of characterizing the behavior in terms of the trait were highly correlated with the probability of false recognition in all experimental conditions. An analysis of the homogeneity of these correlations showed that they were highly homogeneous: $\chi^2(5) = .65$, $p = .99$, for correlations of person judgments with false recognition, and $\chi^2(5) = .81$, $p = .98$, for correlations of behavior judgments with false recognition. Aggregating across experimental conditions, the correlation between the mean false recognition of implied traits and person judgments reached .77, $p < .001$. The corresponding correlation for behavior judgments and false recognition was .59, $p < .001$.

In every experimental condition, the correlation of false recognition with person judgments was higher than the correlation with behavior judgments. Because person and behavior judgments were highly correlated, $r(20) = .80$, $p < .01$, partial correlation analyses are more informative. Although the correlation between person judgments and false recognition (.77) was reduced after controlling for behavior judgments (.61), it remained high and significant ($p < .05$). In contrast, the correlation between behavior judgments and false recognition (.59, $p < .05$) reversed in sign and approached zero (−.06, n.s.) after controlling for person judgments. Analyses within each experimental condition showed the same pattern.

Finally, we regressed false recognition of actor-implied traits on the person judgments, the response times for the person judgments, the behavior judgments, the response times for the behavior judgments (see also footnote 4), and the frequency and valence of the traits. The only significant predictor of false recognition was the person judgments, $t(13) = 2.52$, $p < .026$; for all other predictors, $ps > .34$. The regression analysis accounted for 66.9% of the variance of the false recognition of actor-implied traits. In fact, the person judgments alone accounted for 59.2% of the variance. Thus, the addition of the other five predictors accounted for less than 8% of the variance, and this increment in the variance was not significant, $F < 1$.

Discussion

Explicit person inferences, but not behavior inferences, predicted STIs about the actor, as measured through false recognition. This relation was not qualified by the strength of behavioral inferences. At the same time, the correlation between explicit behavior inferences and false recognition practically dropped to zero after the effect of person inferences was controlled for. These findings suggest that the false recognition reflects spontaneous on-line person inferences rather than merely on-line behavior inferences and associations.

Further, the strength of the relations between explicit on-line person judgments and the false recognition of spontaneously inferred traits remained the same across the different experimental conditions: different rates of presentations of face-behavior pairs (2 s, 5 s, self-pace), size of face–behavior pairs set (60, 120), secondary task, and shallow processing of the behavioral information. These findings suggest that the unintentional person inferences were highly efficient.

[5] The correlation between the false recognition rate for control traits and the judgments did not reach significance in any of the experimental conditions. The range of correlations with person judgments was −.03 to .35. This range was −.08 to .09 for behavior judgments.

Regression analyses showed that the strong relation between the probability of explicitly attributing traits to actors and the false recognition of the STIs could not be accounted for by additional variables such as the valence or the frequency of trait words. The only significant predictor of the false recognition of implied traits was the probability of attributing the trait to the actor.

The present findings extend those of Todorov and Uleman (2002). They found the same relations between explicit trait judgments of the actor and false recognition of implied traits. These relations are somewhat stronger in the current study, most likely because the procedures for measuring the actor judgments were more similar to the procedures used to assess STIs. Most importantly, the pattern was the same in both studies, although the analyses were conducted on completely different sets of face–behavior pairs and responses from different participants.

General discussion

If participants associate an inferred trait with the actor, they should be more likely to falsely recognize this trait in the context of the actor than in a different context. That was the case across the three experiments reported here and across the six experiments reported in Todorov and Uleman (2002). In these experiments, participants were more likely to falsely recognize implied traits paired with actors' faces than implied traits randomly paired with familiar faces.

More important, the current studies showed that the process of binding STIs to actors is relatively independent of attentional resources. All three experiments used manipulations that depleted such resources and interfered with the strategic or controlled processing of the information. Although these manipulations—fast-pace of stimuli, shallow processing, and cognitive load—reduced the overall recognition accuracy of participants, they did not prevent binding STIs to actors' representations. Even when participants saw each face and behavior for only 2 s, when they counted the nouns in each sentence, and when they rehearsed 6-digit numbers during the presentation of faces and behaviors, participants were more likely to falsely recognize actor implied than control traits. In fact, the only manipulation that reduced this effect was the shallow processing of the information while counting nouns. However, even in this condition the effect size was comparable to the effect sizes in the other conditions. The effect sizes across all five experimental conditions were substantial (mean $r = .81$) and highly homogeneous, $\chi^2(4) = 1.94$, $p = .75$.

Finally, Experiment 4 showed that explicit on-line person judgments, but not on-line behavior judgments, predicted false recognition of actor-implied traits, and that this relation was unaffected by different processing conditions. These findings suggest that these highly efficient spontaneous trait inferences were about the persons, not merely the behaviors.

The automaticity of actor-linked STIs

The findings from the false recognition paradigm suggest that the process of binding unintentional inferences to persons' representation is automatic. The basic finding of higher false recognition of implied traits in the context of the actor generalizes across set-size of stimuli, processing time, and processing instructions. The effect was obtained for 36 (Todorov & Uleman, 2002, Experiments 1–4), 60, and 120 (Todorov & Uleman, 2002, Experiments 5 & 6) face–behavior pairs. The effect was obtained for a 2-s presentation time per a face–behavior pair, and for 5 s, 10 s (Todorov & Uleman, 2002, Experiments 1–4), and self-paced presentations. The effect was obtained when participants memorized the information or processed the grammatical parts of the sentences. Finally, the effect was obtained when participants were engaged in a secondary task. These experiments show that the inference process is highly *efficient* and relatively independent of set-size of stimuli, processing time, and attentional resources.

Efficiency is only one dimension of automaticity. Other relevant dimensions include the intentionality of the process, the person's awareness of the process, and the controllability of the process (Bargh, 1994). As pointed out by Uleman (1999), intention is the one criterion that clearly distinguishes spontaneous inferences from controlled processes. STIs are by definition unintentional. Experiment 2, in which participants had a noun-counting goal that precluded even automatic impression formation goals, provides the strongest evidence that STIs in the false recognition paradigm are unintentional.

The current experiments did not directly address the other two dimensions of automaticity—awareness and controllability of the process. However, previous studies have provided evidence that spontaneous trait inferences can occur outside of one's awareness (Lupfer et al., 1990; Moskowitz & Roman, 1992; Newman & Uleman, 1990; Winter & Uleman, 1984). Further, as outlined in Todorov and Uleman (2002), a distinctive advantage of the false recognition paradigm is that the demands of the recognition task and making the inference act in opposite directions. Awareness of making inferences should reduce the likelihood of false recognition.

The controllability of linking STIs to actors was assessed by Uleman and Blader (2002; see also Uleman, Blader, & Todorov, in press). They used the process dissociation procedure introduced by Jacoby (1991; Jacoby, Toth, & Yonelinas, 1993; Jacoby, Yonelinas, & Jennings, 1997), to study the controlled and automatic components of the process of binding of STIs to actors' representations. In the process dissociation procedure, participants are asked to either include or exclude the influence of prior information on their judgments. The difference between these conditions provides a measure of cognitive control. Once control is estimated, the model allows estimation of the contribution of automatic processes to the observed effects. Uleman and Blader modified the procedures of Carlston and Skowronski's (1994; Carlston et al., 1995) savings paradigm. On some trials, participants were told that the behavioral information was randomly paired with the people on the photos, and that they should not use this information in trait judgments of these people. On other trials, participants were told that the information was about the people in the photos, and that they should use it. Uleman and Blader (2002; see also Uleman, Blader, & Todorov, in press) found that controlled processes were practically eliminated by introducing a two-day delay between the study and judgment tasks, but that automatic processes continued to affect trait ratings.

The findings from Experiments 1, 2, and 3, as well as the findings from the above studies, suggest that much of the process of binding STIs to actors' representations can be characterized as automatic.

The nature of STIs

Two experimental paradigms—savings and false recognition—have provided strong evidence that spontaneous trait inferences are bound to the actor's representation. However, as discussed earlier, the nature of these inferences has been ambiguous. It is possible that people make semantic inferences about the meaning of the behavior (e.g., Park, 1989) that get linked to the actor, but do not make trait inferences about the actor. For example, upon observing an honest behavior, people may infer merely that the actor performed an honest act rather than that the actor is an honest person. In Trope's words (Trope, 1986; Trope & Alfieri, 1997), people may not spontaneously go beyond the identification of the behavior.

The findings of Experiment 4 and the correlational findings reported in Todorov and Uleman (2002) provide the first clear evidence that the spontaneous inferences bound to the actor's representation are trait inferences about the actor. Explicit trait judgments of the actor were a strong predictor of the false recognition of implied traits in the context of the actor's face. Explicit behavior judgments were not, when actor judgments were partialled out. Moreover, the various experimental conditions did not affect this pattern. Even when participants had only 2 s to look at the face and read the behavioral information, when they were engaged in a secondary task, and when they were presented with 120 new faces, the explicit on-line actor judgments predicted STIs. In fact, the actor judgments accounted for almost 60% of the variance of the false recognition of implied traits across experimental conditions. In contrast, the behavior judgments were practically unrelated to the false recognition of implied traits after controlling for the effect of person judgments.

Conclusions

Even if we do not intend to make a person-judgment or to form a person impression, but simply attend to another person's behavior, this is sufficient to trigger inferences about the person. These inferences occur even if the behavior presentation is very brief, if we do not attend carefully to the meaning of the behavior, and if our attention is severely constrained by a secondary task. These inferences are about the person and have specific implications for our representation of that person.

References

Bargh, J. A. (1990). Auto-motives: Preconscious determinants of social interaction. In E. T. Higgins, & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 93–130). New York: Guilford.

Bargh, J. A. (1994). The four horsemen of automaticity: Intention, awareness, efficiency, and control as separate issues. In R. S. Wyer, & T. K. Srull (Eds.), *Handbook of social cognition: Vol. 1, Basic processes* (2nd ed., pp. 1–40). Hillsdale, NJ: Erlbaum.

Bargh, J. A., & Barndollar, K. (1996). Automaticity in action: The unconscious as repository of chronic motives and goals. In P. M. Gollwitzer, & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 457–481). New York: Guilford.

Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology, 49*, 1129–1146.

Bassili, J. N. (1989). Traits as action categories versus traits as person attributes in social cognition. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 61–89). Hillsdale, NJ: Erlbaum.

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.

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*, 671–684.

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General, 104*, 269–294.

Francis, W. N., & Kucera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar*. Boston: Houghton Miffin Company.

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin, 117*, 21–38.

Gilbert, D. T., & Osborne, R. E. (1989). Thinking backward: Some curable and incurable consequences of cognitive busyness. *Journal of Personality and Social Psychology, 57*, 940–949.

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.

Lupfer, M. B., Clark, L. F., & Hutcherson, H. W. (1990). Impact of context on spontaneous trait and situational attributions. *Journal of Personality and Social Psychology, 58*, 239–249.

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.

Park, B. (1989). Trait attributes as on-line organizers in person impressions. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 39–59). Hillsdale, NJ: Erlbaum.

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.

Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review, 93*, 239–257.

Trope, Y., & Alfieri, T. (1997). Effortfulness and flexibility of dispositional judgment processes. *Journal of Personality and Social Psychology, 73*, 662–674.

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.

Uleman, J. S., & Blader, S. (2002). A process dissociation analysis of using spontaneous trait inferences. New York University, Unpublished data.

Uleman, J. S., Blader S., & Todorov, A. (in press). Implicit impressions. In: R. Hassin, J.S. Uleman, J.A. Bargh, (Eds.), *The new unconscious*. Oxford University Press, New York.

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.

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, Also see correction in Journal of Personality and Social Psychology (1986), 50, 355.

Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology, 49*, 904–917, Also see correction in Journal of Personality and Social Psychology, (1986) 50, 381.

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