Journal of Experimental Social Psychology 47 (2011) 321–332

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

What will Phil do next?☆

Spontaneously inferred traits influence predictions of behavior

Randy J. McCarthy *, John J. Skowronski *

Department of Psychology, Northern Illinois University, DeKalb, IL, 60115, USA

ARTICLE INFO

Article history:
Received 2 December 2009
Revised 19 September 2010
Available online 21 October 2010

Keywords:
Social judgment
Social inferences
Spontaneous trait inferences
Process dissociation procedure
Automatic processing

ABSTRACT

Past research has shown that perceivers intentionally may make trait inferences about others and use this information to make predictions about these others' future behaviors. Other research has also shown that people can make trait inferences without intent—that is, spontaneously. However, one unexplored avenue is whether spontaneous trait inferences (STI), affect how perceivers predict others' will behavior. Three studies explored this issue. Results from Studies 1 and 2 showed that: (1) exposure to trait-implicative behaviors describing an actor influences subsequent behavior predictions made about the actor in a trait-consistent manner, and (2) predictions occurred regardless of behavior recall, implying that the behavior predictions were derived from prior trait inferences and not from behavior recall. Results from Study 3 bolstered this conclusion by showing that behavior predictions were similar regardless of whether subjects were explicitly instructed to make inferences or not, but that a manipulation known to interfere with inference generation (lie detection instructions) muted behavior predictions. Results from Study 3 also suggested that the prediction effects had both automatic and controlled components, and that reductions observed in the lie detection condition of Study 3 were caused by alterations in the automatic influence of trait knowledge to the behavior predictions. These results suggest that STI may be causal inferences about the actors' dispositions.

© 2010 Elsevier Inc. All rights reserved.

Humans typically live in highly social environments. Accordingly, the ability to understand other people and the implications of their behavior is seemingly a critical skill for successfully navigating such environments. One potential indication of the benefits of social inference-making is its ubiquity: Research strongly suggests that perceivers routinely and spontaneously make trait inferences about an actor after observing or learning about an actor's trait-implicative behavior. For example, on discovering that Amanda aced her statistics exam, people may infer that she is intelligent. This occurs even if people do not intend to infer trait information and can even occur without awareness that such an inference has occurred (for reviews see Skowronski, Carlston, & Hartnett, 2008; Uleman, Newman, & Moskowitz, 1996; Uleman, Saribay, & Gonzalez, 2008). While past research has focused on detecting such inferences, the present research focuses on one of the proposed functions of such inferences, to inform our predictions of others' behavior.

We contend that the ability to make effortless inferences is a critical social skill; people are not social perceivers merely taking in information from their environments, but people are immersed within, and thus interact with, their environments. How can such inference-making be functional? One example derives from the idea that people often want to feel as if they can predict and control their worlds (e.g., Biner, Angle, Park, Mellinger, & Barber, 1995; Bruner, 1957). Social inference-making can give a perceiver the desirable feeling that they know what to expect from others. More concrete functional benefits can be derived from the idea that the ability to use inferences to accurately make predictions about others' behaviors allows the formation of reasonable action plans that maximize the outcomes that one might obtain from interactions with others. Indeed, research has shown that failure to accurately interpret others' behaviors may leave one vulnerable to exploitation or unable to cooperate on joint tasks (see Komorita & Parks, 1995 for a review of interpersonal relations strategies in social dilemmas). Additional evidence pertaining to functionality comes from observations of individuals who often have difficulty understanding others' internal worlds, such as those with autism (see Chakrabarti & Baron-Cohen, 2008; Klein, Cosmides, Murray, & Tooby, 2004). Such deficits are often detrimental to autistics' ability to navigate their social environments. In fact, the difficulties exhibited by the inference-impaired has led to speculations that the ability to accurately detect others' traits might have its basis in evolution (Nesse, 2005; Schaller, 2008; Schaller, Park, & Kenrick, 2007; Schaller, Simpson, & Kenrick, 2006).

One idea that is crucial to the functionality of trait inferences is the perception that traits are characteristics that are internal to actors and

☆ Some studies described in this article were conducted as part of a Master's Thesis conducted by the first author under the supervision of the second author.

* Corresponding authors.
E-mail addresses: rmccarthy3@niu.edu (R.J. McCarthy), jskowron@niu.edu (J.J. Skowronski).

0022-1031/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jesp.2010.10.015

causal to behavior (e.g., he donated money *because* he is generous). This causality assumption is a cornerstone of many theories of dispositional inference (e.g. the Naïve Scientist, Heider, 1958; Correspondent Inference Theory, Jones & Davis, 1965; Covariation Theory, Kelley, 1967). That such an assumption exists is supported by data showing that actor traits seem to be treated as actor-internal forces that act across situations (Newman, 1996; Pittman & Heller, 1987; also see Shoda & Mischel, 1993). This notion is further supported by research showing that people use their prior inferences to make predictions about others' future behaviors (Dunning, Griffin, Milojkovic, & Ross, 1990; Griffin, Dunning, & Ross, 1990; Newman, 1996; Nussbaum, Trope, & Liberman, 2003). Some studies (Newman, 1991) even explicitly link the development of inference-making abilities to the construction of behavior predictions. For example, Newman's (1991) data suggest that children who were especially likely to form trait inferences also were especially likely to make trait-consistent behavior predictions.

However, whereas prior research focused on behavioral predictions made consequent to *intentional* inference formation, the present studies assess whether such predictions occur after people *spontaneously* infer traits about others. This issue is theoretically important because one unresolved issue in the STI literature concerns the exact nature of STIs and whether they are the same as, or have the same properties as, intentionally-generated inferences.

For example, consider the data collected by Carlston and Skowronski and their associates (see Skowronski et al., 2008). Results from their studies demonstrate that trait labels derived from trait-implicative behavior descriptions become linked to mental representations of actors. Clearly, for such linkages to form, people must activate a trait label on encountering an actor's behavior and process the label in working memory so that it becomes associated with the actor. However, as Bassili (1989) has noted, a linkage or an association is *not equivalent* to a trait inference. An association merely implies that two constructs have been linked in memory. A trait inference goes beyond this simple linkage, implying that the trait is a property of the actor. For example, upon reading that *Brett found a wallet and returned it to the owner with all of the money still inside*, perceivers may infer that Brett is honest (a trait inference) and would likely perform other honest behaviors. Alternatively, perceivers may simply summarize the behavior with a trait label (e.g., that was an honest behavior) without inferring dispositional information about the actor. The trait concept (e.g., honest) then becomes associatively linked to the actor (e.g., Brett) in memory. Indeed, Carlston, Skowronski, and their associates have conducted an extensive series of studies designed to show that STIs have different psychological properties than actor–trait associations (see Carlston & Skowronski, 2005; Crawford, Skowronski, Stiff & Sherer, 2007).

The present studies speak to this association versus inference debate. If the kinds of spontaneous trait inferences that people make on exposure to behaviors are essentially the same as the inferences that are made when people have explicit inference-making goals (usually instructions given by the experimenter to form an impression of an actor), then one might expect people to use those inferences to make predictions of an actor's future behaviors. Moreover, the kinds of post-inference behavior prediction effects observed by researchers such as Newman (1996) ought to be observed, regardless of whether perceivers have been given an explicit impression formation processing goal or not. Alternatively, as Bassili (1989) might claim, if spontaneous inferences are simply labels that become associatively linked to actors, but have no causal implications, then predictions derived from intentional inference processing may be considerably stronger than predictions derived from spontaneous inferences.

We suspect that the former might be true, and not the latter. One reason lies in results of studies pursuing Spontaneous Causal Inferences (SCIs; Hassin, Bargh, & Uleman, 2002). The idea underlying SCI research is that, on reading causally implicative scenarios, participants spontaneously infer causes for the events described. Tellingly, the methodologies used to study STIs and SCIs overlap considerably. For example, in STI research, when participants read a trait-implicative scenario, the implied trait often serves as a better recall cue than words that were actually presented in the scenario (see Winter & Uleman, 1984). For instance, participants better recall the sentence "*Henk refused to try a new brand of beer*," when cued with the implied trait (e.g., *picky*) than when cued with a word that was in the sentence (e.g., *brand*). Likewise, in SCI research, a causal action serves as an excellent recall cue for cause-implicative scenarios (Hassin et al., 2002). For example, the causal action *pickpocket* is a more effective cue for the scenario "*After spending a day exploring the beautiful sights in the crowded streets of New York, Jane discovered that her wallet was missing*" than the word *sights*, despite the latter word's presence in the scenario. The similarity in the data collected in both lines of inquiry suggests that the same causal implications explored in the SCI research may also characterize STIs.

Given these precedents, we hypothesize that STIs made about actors will subsequently be used to make predictions of future behaviors, and that such predictions will be equivalent to the predictions made when one has explicit impression formation goals. Importantly, we expect that such predictions will be made, *even in the absence of explicit memory* for the behaviors on which these predictions are ultimately based. Such findings would be consistent with the notion that the trait construct extracted from the initial behavior, and not the initial behavior itself, forms the basis of future behavior predictions.

Several pilot studies and three experiments explored these ideas. Because our behavior prediction studies employed trait-implicative stimuli that had not been used previously, we thought it was prudent to first document the emergence of STIs using these new stimuli. This goal was pursued in a series of pilot studies. Using these new stimuli, Studies 1 (actor-behavior matching task) and 2 (behavior likelihood judgment task) were designed to show that exposure to trait-implicative behaviors describing an actor influences subsequent behavior predictions made about the actor. An additional goal of Studies 1 and 2 was to show that prediction effects occurred regardless of behavior recall, implying that predictions were derived from trait inferences and were not directly generated from behaviors. Finally, the prime goals of Study 3 were to show that behavior predictions were similar regardless of whether or not participants were explicitly instructed to make inferences, and that a manipulation known to interfere with inference generation (lie detection instructions) muted such behavior predictions. Study 3 also examined the contributions of controlled processes and automatic processes to these behavior predictions.

Pilot studies

The STI replication study exposed subjects to descriptions of actor behaviors. Some behaviors implied a trait (*trait implied* or TI behaviors); some did not (these were either *trait not implied* or TNI behaviors, or were *neutral* behaviors). The behavior descriptions used in TNI trials minimized a potential confound in past research. In such research, the trait-implicative sentences and non-trait-implicative sentences (equivalent to our neutral sentences; see Carlston & Skowronski, 1994) often differed in a number of uncontrolled ways. These included differences in semantic content and description complexity. These pilot studies eliminated some of these confounds. The TNI descriptions used in these pilot studies exhibited a considerable amount of semantic overlap with the TI descriptions and were designed to be equal to the TI sentences in terms of complexity and ease-of-reading.

Materials development pilot studies

To obtain such behaviors, fifty-one candidate behavior pairs were developed. One version of each pair (e.g., *I got the same grade as my*

group for the class project even though I did not do any work) was intended to imply that the narrator possessed a specific trait (*lazy*). Modifications made to the second version of each behavior attempted to mute the behavior's trait implication but still possess considerable semantic overlap with the TI version (e.g., *I got the same grade as my group for the class project and we all did the same amount of work*).

In the first pilot study, subjects ($N=39$) from two groups read one version of each behavior, with different versions shown to each group. Subjects then rated how well a given trait (one thought to be implied by the TI description) reflected the narrator's disposition. Sixteen pairs with the highest mean difference scores across the two versions of each behavior were selected for subsequent studies. Further, the sixteen TI behaviors chosen for use in the studies all implied unique traits.

In a second pilot study, subjects ($N=20$) read each behavior pair (TI and TNI) and indicated whether one was easier to read than the other. None of the behavior pairs differed in judged readability (all $ps>.50$). Hence, it is not plausible that inferences are less likely to be derived from the TNI sentences because they are more complex than the TI sentences.

STI replication pilot study

Subjects

Ninety students from Northern Illinois University volunteered for the study. They received credit toward completion of a course research option as compensation.

Procedure

Tasks were presented on a computer workstation running Direct RT software (Empirisoft, 2004). The procedure modified that described by Carlston and Skowronski (1994).

In the study's initial *exposure* task, subjects encountered a series of trials. Each trial presented a photo and behavior description (photo above, behavior below). Subjects were (falsely) told that each description was provided by the person in the photo accompanying the description. Subjects were told that the purpose of this exposure was to familiarize them with stimuli to be used in later tasks. Because subjects were told only to familiarize themselves with the materials (and not to form impressions) any inferences are considered spontaneous.

Each photo-behavior pair appeared on the computer screen for nine seconds then disappeared. Subjects simply read each description and looked at each photo. After nine seconds, a blank screen appeared for two seconds and the next photo-behavior pair appeared. This continued until all photo-behavior pairs were viewed.

Subjects viewed one of two different sets of behaviors. The sets (referred to as sets A and B) were structured such that if one group received the TI version of a description, the other group received its TNI complement. Pairings of photos and behaviors were randomized for each subject; hence, results cannot be attributed to idiosyncratic behaviors, traits, photos, or pairings. Additionally, eight neutral behaviors (e.g., *I turned the light on when it got dark*) also appeared in each set. These served as an additional set of control trials that allowed for comparisons to be made to past research using such behaviors. These behaviors were the same for all subjects. Overall, then, each subject viewed 24 photo-behavior pairs (8 TI, 8 TNI, and 8 neutral).

After a short filler task, subjects encountered a *learning* task in which, on each trial, they viewed a person photo accompanied by a word. Subjects were told to try to remember the word paired with each photo and told that they would later be tested on their ability to recall these pairs. Four types of actor-word pairs appeared. For some (*TI trials*), the photo was paired with the trait word implied by the trait-implicative behavior shown with the photo during the exposure task. Other pairs (*TNI trials*) consisted of trait words paired with photos that previously accompanied the TNI version of a behavior. Still other pairs (*neutral trials*) presented photos that were accompanied by neutral behaviors during the exposure task paired with novel trait words (i.e., not implied by any behaviors). Finally, twenty-four previously unseen photos were shown with non-trait words (e.g., tree, car). The purpose of these trials was to provide photos that could serve as familiar control foils in a subsequent recognition/choice task.

After a brief filler task, subjects completed the *recognition/choice* task. In this task, two photos appeared side-by-side in the upper half of the computer screen and a trait word appeared centered below them. Subjects pressed a key ("Z" for the left photo and "/" for the right photo) to indicate their judgment as to which photo was paired with the trait word during the learning task. Each choice was made between a familiar control photo (each previously shown with a non-trait word during the learning task) and a photo shown with a trait word during the learning task. The position of the photos on the screen (left, right) was randomly determined for each subject.

In a final *behavior memory* task, subjects attempted to recall the behavior shown with each photo. On each trial, a photo from the initial exposure task appeared on the screen accompanied by a text box. Subjects were prompted to type in the behavior description, as exactly as possible, that was shown with each photo viewed during the exposure task.

After all trials had been completed, a screen appeared thanking the subjects and indicating that they were done with the study. Subjects were then debriefed and dismissed.

Savings emerged in actor recognition/choice

For each subject, the numbers of correctly chosen photos within each behavior type condition (TI, TNI, neutral) were tabulated. A subject's score within each condition ranged from 0 to 8. Results are described in terms of the percentage of photos correctly chosen; these were examined using a Behavior Type (TI vs. TNI vs. neutral; within-subject) × Counterbalancing Group (A vs. B; between-subject) ANOVA.

The analysis yielded only a behavior type main effect, $F(1,178)=7.85$, $p<.01$, $\eta^2=.081$ (counterbalancing group $F(2,89)=.621$, *ns*; interaction $F(2,178)=.183$, *ns*). To clarify the meaning of this effect, a planned linear contrast was conducted. Weights of +2, −1, and −1 were applied to scores in the TI ($M=95.5\%$ correct, $SD=.09$), TNI ($M=91.8\%$ correct, $SD=.11$), and neutral ($M=91.1\%$ correct, $SD=.12$) conditions, $F(1,178)=4.99$, $p<.01$. These results suggest that reading a trait-implicative behavior elicited inferences about an actor; these inferences then made it especially easy to learn an association between the trait and a photo when the trait was later paired with the actor's photo in the learning task.

A second planned contrast evaluated the extent to which trait inferences in response to TNI behaviors were muted. This was done by examining a contrast in which weights of 0, −1, and +1 were assigned to the TI, TNI, and neutral conditions. The TNI recognition score did not significantly differ from the recognition score for the neutral photos, $F(1,178)=1.26$, *ns*.

Savings did not depend on behavior recall

One concern is that savings effects in the choice data might not come from spontaneous inferences made during exposure, but that they might derive from a process that is mediated by behavior recall during the recognition/choice task. Accordingly, an analysis was performed to examine this issue.

Two trained scorers who were blind to the hypotheses of the study examined the content of subjects' attempts to recall the original behaviors. They assigned each recall attempt to one of three categories: no behavior memory, partial memory; or complete or almost complete memory. A concrete example is illustrative; the recall attempts "*don't know*," "*he missed his haircut*," and "*Forgot to get her hair cut although it was the only thing she had going on that day*" for the behavior "*I did not go to my haircut appointment. I can't believe I missed it; it was the only thing I had to do today*" were coded as no memory, partial memory, and complete or almost complete memory, respectively. Raters agreed 88.75% ($\kappa=.76$) of the time; disagreements were resolved by a third rater.

To measure the influence of recall on recognition/choice performance, subjects' behavior recall performance was entered as a categorical variable in a mixed-model moderated hierarchical regression analysis predicting correct photo-trait matches on a trial-by-trial basis (this trial-by-trial examination accounts for the magnitude of the error *df*s in the analyses reported later). Only the trials that depicted similar behaviors (TI, TNI) were included. Hence, each subject contributed 16 trials (8 TI and 8 TNI) to the analysis.

If subjects generated inferences only after recalling the original behaviors, one would expect that recognition/choice correctness rates would be especially high when behaviors were correctly recalled.[1] This was not the case. The behavior type effect reached significance, $F(2,1362) = 8.48$, $p < .004$; however, the Behavior Type (TI vs. TNI) × Recall (no memory vs. partial memory vs. complete memory) interaction was not statistically reliable, $F(2,1360) = .51$, *ns*. The analysis also did not yield a main effect for recall, $F(2,1362) = 1.88$, *ns*. Separate examination of the behavior recall-correct choice/recognition relation only for trait-implied behaviors also yielded null effects for the recall-correct choice relation, $F(2,635) = 1.61$, *ns*.

These results converge on a common conclusion: Correct recognition of actors paired as with traits was *not* dependent on behavior recall. Thus, these data do not suggest that subjects generated trait inferences after recalling behaviors paired with the photos. Instead, the results suggest the heightened recognition/choice performance observed in the TI condition reflects spontaneous trait inferences made at behavior encoding about actors, conceptually replicating conclusions from earlier studies (Carlston & Skowronski, 1994; Todorov & Uleman, 2002).

However, we note that the results of this pilot study go beyond a simple replication of these earlier results. First, our results show that evidence of spontaneous inferences obtained by means of the free recall savings task used in earlier work (e.g., Carlston & Skowronski, 1994) can also be obtained by means of the photo choice task that we employed in the current pilot study. While both tasks assess the same theoretical person–trait association, one might argue that the photo choice task might be preferable because it can be objectively scored and does not require the kinds of coder judgments (e.g., to count synonyms or misspelled words as correct or incorrect) required by the standard savings task. However, the choice task may not be perfect—the high correctness rates obtained might cause researchers to be concerned about ceiling effects. Nonetheless, the fact that similar effects emerge across different dependent measures adds convergent validity to the result obtained in the current study.

A second point to be made about this pilot study is that, other than their trait implications, the TI and TNI behaviors used were similar in many ways. Despite these similarities, these behavior types yielded differences in the apparent STI activity levels prompted by such behaviors. Such results suggest that the data yielded by earlier studies are not confounded by non-trait differences between trait-implicative behaviors and non-trait-implicative behaviors (equivalent to our neutral behaviors) used in earlier studies. In fact, given these results, we decided not to use the neutral behaviors in Studies 1 and 2.

Study 1

Overview

While the pilot study (and past research) focused on detecting trait inferences, Study 1 explored whether these inferences affect predictions of future actor behavior. During Study 1, participants were first exposed to photo-behavior dyads; as evinced in the pilot study, pairing a trait-implicative behavior with an actor should elicit a trait inference. In a second task subjects tried to match the actors with whom the behaviors were paired to new behaviors the actors were thought to be likely to perform. Therefore, if the trait inferences are about the actors' dispositions (i.e., trait inferences) then participants should match the actors to new behaviors that are consistent with the inferred traits. In addition to documenting the utility of inferences (e.g., that they can influence behavioral predictions), such results would also provide evidence supporting the proposition that STI are about the actors' characters and have causal implications.

Method

Stimuli development and selection

Study 1 required a set of new behaviors that would serve as the targets of subjects' predictions. Accordingly, candidate behaviors were developed by the article authors. Then, in a pilot study, subjects ($N = 40$) rated each of these new behaviors on the trait implied by the TI behaviors validated in the pilot study. The new behavior evincing a trait rating that yielded the smallest difference from the score evinced by the trait-implicative behavior used in the pilot study (a value known from one of the pretests described earlier) was selected. Thus, each trait-implicative behavior used in Study 1's exposure phase and its corresponding new behavior both implied the same trait to approximately the same degree.

To select the photos used in Study 1, subjects in a pilot study ($N = 28$) rated forty-eight photos on all sixteen of the critical traits used in Study 1. The sixteen photos most neutral on all sixteen traits were chosen for use. This was done to minimize the extent to which behavior predictions would be affected by actor appearance.

Subjects

Sixty-four students from Northern Illinois University volunteered for the study and were compensated with credit toward completion of a course research option. Data from one subject was excluded from analyses: Examination of that subject's responses indicated that the subject failed to follow task instructions.

Procedure

As in the pilot study, subjects were told that the purpose of the initial exposure task was to familiarize them with the types of stimuli that would be used later in the study. Hence, inferences made about the actors in this task were made spontaneously. The exposure task was identical to that used in the pilot study, except that the neutral behaviors were not used in the current study. Thus, subjects were exposed to 16 actor photos, each paired with a behavior. For half of the subjects, half of the behaviors were in the TI form and the other half were in the TNI form. The other subjects viewed the complementary form of the behaviors. A variable reflecting these two behaviors sets (sets A and B) was included in analyses.

Subjects then completed a short filler task (~5 min). Afterwards, an experimenter entered the room, gave subject two decks of cards—one deck contained actor photos and one deck contained new behaviors—and an answer sheet, and read a script explaining the next task. The task was to match the actor in the photo with a new behavior that they would be likely to perform. Subjects were to do this by physically pairing photo cards with behavior cards by laying the cards out on the table and arranging the cards into pairs, then recording each pair (numbers on the photos, letters on the behaviors) on the answer sheet.

After recording their answers, subjects completed the same photo-cued behavior recall task used in the pilot study. Subjects were then debriefed and released.

[1] The dependent measure in the regression analysis is dichotomous; hence, some might argue that these data are best analyzed using logistic regression. However, in this case we were concerned about comparability between this regression analysis and the ANOVA that was performed on the proportion of correct matches. In any case, the exact analysis conducted does not matter to the conclusion that memory does not seem to play a role in these behavior prediction results.

Results and discussion

Behavior matching was affected by prior exposure to trait-implicative behaviors

The main dependent variable was the frequency with which actors were assigned a new behavior that matched the trait implied by the behavior shown with the actor during the exposure task. On TNI trials an assigned behavior was scored as a match if the predicted behavior was consistent with the trait implications of the TI form of the behavior. There could be from 0 to 8 matches for each behavior type.

A 2 (Counterbalancing Group: A vs. B; between-subjects) × 2 (Behavior Type: TI vs. TNI; within-subjects) mixed-model ANOVA yielded only a significant behavior type main effect, $F(1,61) = 20.095$, $p < .01$, $\eta^2 = .248$ (interaction $F(1,61) = 1.35$, *ns*, counterbalancing group main effect, $F(1,61) = .00$, *ns*).[2] The means for this significant effect show that subjects more often matched new behaviors to informants who previously engaged in behaviors implying the same traits as the new behaviors ($M = 2.54$, $SD = 2.23$) than they did to informants who previously conveyed behavior descriptions that did not imply those traits ($M = 1.49$, $SD = 1.60$). This result suggests that during the initial exposure task subjects inferred dispositions about actors, and that during the matching task these dispositional inferences affected the likelihood of matching actors to future behaviors consistent with these inferences.

New behavior predictions were not related to old behavior recall

A subsidiary analysis investigated the possibility that subjects' matching performance was caused by recall of the behaviors presented in the exposure task. Prior to analysis, two raters scored the subjects' recall performance using the same system described in the pilot study. There was 93.1% ($\kappa = .91$) agreement between the raters; all discrepancies were resolved by a third rater.

The behavior recall scores were then entered as a predictor (with behavior type) in a mixed-model moderated hierarchical regression analysis predicting matching performance. As in the pilot study, results from each individual trial were entered into the analysis. The analysis yielded a behavior type effect, $F(1,942) = 21.83$, $p < .0001$. The Behavior Type (TI vs. TNI) × Recall (no memory, partial memory, and complete memory) interaction was not statistically reliable, $F(2,940) = .34$, *ns*. The main effect for subjects' recall performance similarly yielded null results, $F(2,942) = .37$, *ns*. Finally, the extent to which behavior recall predicted correct choice/recognition only for trait-implied behaviors also yielded null effects, $F(2,439) = .72$, *ns*.

Collectively, these results suggest that subjects' matching performance cannot be accounted for by their recall for the original behaviors. Instead, the results suggest that prior inferences about actors to influenced predictions about the actors' future behaviors.

Study 2

Overview

One might argue that the spontaneous inferences that people make about others are largely evaluative, a proposition supported by the finding that affective reactions to stimuli seem to be ubiquitous (e.g., Fazio, 2001; Zajonc, 1984). One might also argue that such spontaneous evaluations might lead people to make behavioral predictions that are evaluatively consistent with the original behavior. However, Study 1's results argue against this evaluative inference interpretation. Recall that in Study 1, a subject had a choice among behaviors that could be paired with an actor photo. Many of these behaviors were evaluatively consistent with the original behavior performed by an actor, but only one was consistent with the *exact* trait implications of that behavior. The spontaneous evaluation-only idea cannot easily account for the fact that people's choice of future actor behaviors went beyond evaluative fit and matched the exact trait implications of the actor's prior behavior.

Nonetheless, the issue of whether inferences were trait inferences or evaluative inferences is important, hence, it was examined more closely in Study 2. Subjects in this study directly judged the likelihood that a previously encountered actor would perform each of several new behaviors. Because the exact content of the predicted behaviors was systematically manipulated, this method allows conclusions to be drawn about the exact nature (evaluative or trait) of the spontaneous inferences drawn on exposure to an actor's initial behavior.

If spontaneous inferences are solely evaluative, then predictions of future behaviors should be strongly affected by the behaviors' evaluative content, regardless of the trait implications of a behavior. That is, if Jessica's behavior prompts a perceiver to perceive Jessica positively, the judged likelihood that she would help the homeless (*generous*) should be increased relative to a base rate just about as much as the judged likelihood that she would refuse an offered bribe (*honest*). However, if the original behavior prompted a trait inference, the judged probability of a future behavior should be especially high if the trait implications of the new behavior match the trait implications of the original behavior. Hence, if an original behavior prompts a perceiver to see Jessica as honest, the judged probability that she would refuse an offered bribe should be higher than the judged probability that she would help the homeless.

Certainly, even if initial inferences involve traits, one would certainly expect the judged likelihoods of evaluatively consistent behaviors to be heightened, and the likelihood of evaluatively inconsistent behaviors to be lowered: Trait judgments typically have evaluative implications, and these should be reflected in *halo effects* in the likelihood judgments (see Carlston & Skowronski, 2005). However, the spontaneous trait inference hypothesis predicts a *specificity effect* in the data, such that judged future behavior likelihoods should be especially high when the trait implications of behavior match the trait implications of the original behavior.

Method

Subjects

Sixty-three students from Northern Illinois University volunteered for the study and were compensated with credit toward completion of a course research option.

Procedure

As in the previous studies, in an exposure task subjects encountered actor photos paired with behaviors (either TI or TNI) supposedly performed by those actors. The sixteen photos and behaviors were exactly the same as those used in Study 1, with the behavior sets counterbalanced across two subject groups (a TI behavior in one group was a TNI behavior for the other). Subjects viewed these materials after receiving the same familiarization instruction used in the previous studies. After completion of this task, subjects completed a short filler task (~5 min).

Next, subjects completed a behavior prediction task. In this task, each photo was shown three times, and each time the photo was paired with a new behavior (the new behaviors developed for Study 1). These photo-behavior pairs appeared one-at-a-time; presentation order was randomly determined for each subject. After encountering each pair, subjects rated the likelihood that the actor would perform the behavior. Ratings were made on an ordered category scale (1 = *not at all likely*, 7 = *very likely*). Note that each

[2] Carlston and Skowronski (2005) have contended that affect (via the negativity bias) might be a signature of trait inferences. However, because of the general tendency for positive events to be seen as more likely than negative events, the same negativity bias may not occur for predictions of future behavior. Indeed, valence was included as a variable in the design of Studies 1–3, but the results were not clear. To save space and reduce complexity, we omit discussion of these results.

behavior appeared three times, each with a different actor: Once as a behavior that *matched* implications of the TI behavior used in the exposure task, once as a behavior that was *evaluatively consistent* with a TI behavior (but that reflected another TI trait), and once as a behavior that was *evaluatively inconsistent* with a TI behavior (but that reflected another TI trait).

After completing the behavior prediction task, as in the previous studies, subjects completed a cued recall task in which they were given a photo from the exposure task and were asked to report the original behavior with which that photo was paired. On completion of this task, the subjects were debriefed and released.

Results and discussion

Exposure to trait-implicative behaviors affected behavior predictions

Subjects' ratings were analyzed using a 2 (Behavior Type: TI vs. TNI; within-subject) × 3 (New Behavior: trait match vs. same valence vs. opposite valence; within-subject) × 2 (Stimulus Group: A vs. B; between-subject) mixed-model ANOVA. This analysis yielded a significant Behavior Type × New Behavior interaction, $F(2,122) = 17.771$, $p < .01$, $\eta^2 = .226$ (see Fig. 1).

The interaction's meaning was explored in two ways. First, a series of *t* tests compared likelihood ratings for photos shown with TI behaviors with those shown with the TNI behaviors on trait match, valence-consistent, and valence-inconsistent new behaviors. To control for study-wise error rates, an α level of $p = .0167$ was adopted (Bonferroni adjustment). Actors shown with TI behaviors ($M = 5.01$, $SD = .94$) were rated as more likely to engage in new trait match behaviors than those shown with TNI behaviors ($M = 4.08$, $SD = .85$), $t(62) = 6.258$, $p < .01$, Cohen's $d = 1.04$. Similarly, actors were rated as more likely to perform new behaviors that were evaluatively consistent with the target trait if actors were previously shown with TI behaviors ($M = 4.22$, $SD = .82$) instead of TNI behaviors ($M = 3.70$, $SD = .80$), $t(62) = 4.007$, $p < .01$, Cohen's $d = .65$. Finally, actors were seen as less likely to engage in new evaluatively inconsistent behaviors if the actors were previously shown with TI behaviors ($M = 3.39$, $SD = .95$) instead of TNI behaviors ($M = 3.62$, $SD = .70$), $t(62) = 1.95$, $p = .056$, Cohen's $d = .28$.

The meaning of the interaction was also probed by examining the likelihood ratings within each behavior type. These were also conducted by means of Bonferroni-adjusted pairwise *t* tests ($\alpha = .0167$). Results revealed that higher likelihoods were assigned to behaviors that matched trait implications of the original behavior than to behaviors that merely matched the original behavior's valence, both for the TI behaviors ($t(62) = 7.642$, $p < .01$, Cohen's $d = .96$), and the TNI behaviors ($t(62) = 3.60$, $p < .01$, Cohen's $d = .45$). Hence, predictions cannot be a sole function of spontaneous evaluative impressions—if they were, prediction extremity would not vary across these two behavior types. Nonetheless, evaluation did matter: For the TI behaviors, likelihood judgments were greater for valence-consistent than for valence-inconsistent behavior ($t(62) = 4.809$, $p < .01$, Cohen's $d = .61$). The results also suggested that actors performing TNI behaviors were not judged to be more likely or less likely to perform evaluatively consistent behaviors or evaluatively inconsistent behaviors, ($t(62) = .558$, *ns*).

When considered in their entirety, these results suggest that spontaneous trait inferences were made about actors, and that these were more extreme for actors who performed TI behaviors than those who performed TNI behaviors. The results also suggest that the inferences were used to make predictions about subsequent behaviors to be performed by the actors.

New behavior likelihood judgments were not related to old behavior recall

Subsidiary analyses investigated the idea that subjects' ratings were not a reflection of spontaneous trait inferences, but were caused by recall of behaviors presented in the exposure task. Prior to analysis, two raters coded the subjects' recall data using the scoring system described in the pilot study. There were so few behaviors that were classified as reflecting complete recall that, for purposes of analysis, the complete recall category was merged with the partial recall category. A 2 (Behavior Type: TI vs. TNI) × 2 (Group: A vs. B) × 2 (Recall: no memory vs. at least some memory) mixed-model hierarchical regression examined whether recall was related to likelihood judgments reported for the new behaviors that matched trait implications of the original behavior.

The analysis did not yield a significant Behavior Type × Recall interaction, $F(1,942) = .23$, *ns*, the effect that would be expected if behavior recall was needed to generate predictions of future behavior. However, the analysis yielded a relation between recall and the likelihood ratings given to the new behaviors, $F(1,942) = 5.50$, $p < .05$, $\eta^2 = .14$. The means for this effect show that likelihood ratings were especially high when there was partial memory for the original behavior, regardless of behavior type (TI: no memory $M = 4.98$; partial memory $M = 6.361$ TNI: no memory $M = 4.06$; partial memory $M = 4.97$).[3] While this effect might seem to be suggestive of a role for memory in judgments of future behavior, that conclusion is illusory. Because of the large numbers of behaviors viewed in Study 3, behavior cued recall rates were quite poor. Only 20 behaviors (10 TI and 10 TNI) were classified as showing even partial recall in the actor-cued behavior recall task. This low recall rate makes it unlikely that behavior recall fully accounts for the likelihood data collected in Study 3. Indeed, a regression analysis that discarded correctly recalled trials yielded a robust behavior type main effect, $F(1,924) = 46.74$, $p < .0001$.

Collectively, the results of the analyses that examine the relation between behavior recall and the likelihood judgments suggest that subjects' judgment about the likelihood of engaging in future behaviors cannot be accounted for by their recall for the original behaviors. Instead, the results are consistent with the idea that people used prior inferences about actors to generate predictions about their future behaviors.

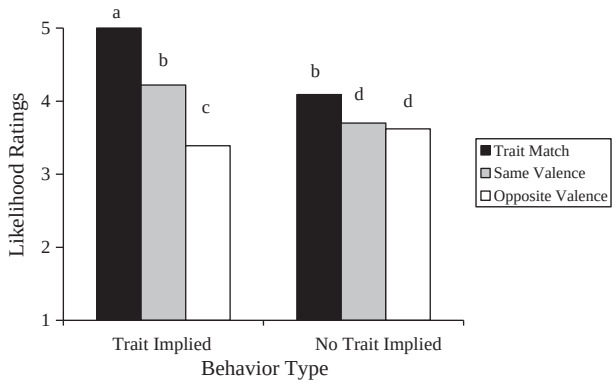


Fig. 1. *Note.* Mean new behavior likelihood ratings from Study 2; different letters indicate that means are significantly different, $p < .05$.

Study 3

Study 2 examined whether spontaneous inferences were solely evaluative or whether such inferences were trait specific. Although subjects rated the actors as likely to behave in ways that were evaluatively consistent with their original behaviors, and unlikely to behave ways that were evaluatively inconsistent with their original behaviors (a *halo effect*), ratings were especially high when the

[3] These means are least-squares means reflecting the effects as they are adjusted for the other terms entered into the regression model.

predicted behaviors matched the exact trait exemplified by the original behaviors. This *specificity effect* in the pattern of data obtained is difficult to reconcile with the evaluative-inference-only idea.

In this regard, we note that the Study 1 data had similar theoretical implications. However, the method used in Study 1 allowed a dependency to exist in the dependent measure: A behavior that was predicted for one actor could not be predicted for a second actor. While subject choices for an actor's future behaviors in Study 1 seemingly went beyond evaluative matches between past and future behaviors, reflecting trait-based matches, one could argue that this result might have been an artifact of the inter-behavior dependencies in the choice/matching measure. The results of Study 2 render this interpretation less likely. Instead, results from both studies converge on the idea that encountering a description of an actor's original behavior spontaneously prompted a trait inference; those inferences then later affected predictions of the actors' future behaviors.

However, while this conclusion is collectively suggested by the results that we have so far described, those results are not conclusive. That is, the results of the pilot study suggest that exposure to trait-implicative behaviors prompts STIs; the results of Studies 1 and 2 suggest that exposure to trait-implicative behaviors prompts behavior-consistent predictions. While the data from all 3 studies collectively suggest that these effects are not caused by recall of the original trait-implicative behaviors, they do not directly show that the predictions observed in Studies 1 and 2 were *caused by* the inferences that were documented in the pilot study.

One way to do so would be to be to conduct a study that assessed both STIs and behavior predictions and show via mediational analyses that the STIs mediate behavior predictions. However, as deceptively simple and attractive as this strategy might be, it poses a significant technical problem: Any method used to assess whether inferences are made might itself influence the behavior predictions that would be made at a later time.

Instead, in Study 3 we chose an alternative strategy. Prior to encountering the trait-implicative behaviors, some subjects in Study 3 were *explicitly instructed* to form an impression of each actor. We compared the behavior predictions made by these explicitly instructed subjects to the predictions made by subjects given the same familiarization instructions used in the previous studies and that presumably made STIs. We reasoned that if uninstructed subjects spontaneously formed the same kind of trait inferences that were made by explicitly instructed subjects, and used them in the same way, then the behavior predictions made in the two conditions should not differ. Differences in the behavior predictions made would suggest differences in inference strength or the nature of the inferences prompted in the two conditions.

We also thought that the case for behavior predictions being derived from STIs would be further enhanced if we could demonstrate that behavior predictions were inhibited by task instructions known to impair STI formation. Accordingly, a third group of subjects in Study 3 tried to discern whether the trait-implicative behaviors that they encountered were each a lie or were true. This latter *lie detection* instruction has been shown to interfere with STI formation (Crawford et al., 2007). Thus, if behavior predictions depend on STIs, and if the lie detection instructions interfere with STIs, then subjects in the lie detection conditions should make less extreme behavior predictions than those in the familiarization and explicit inference conditions.

One final issue addressed in Study 3 was the extent to which the application of STIs to behavior predictions was automatic, and the extent to which the expression of STI were controllable. To do so, in Study 3 the same behavior prediction ratings that were made in Study 2 were made in one of two instruction conditions. Some were made while subjects were told to try and remember the behavior that was shown with the actor earlier in the experiment *in order to help make their prediction*. Other ratings were made while subjects tried to remember the behavior shown with the actor *so as to try and not use this trait information while making their prediction*. Therefore, if a process expresses itself (e.g., a trait inference influencing a behavior prediction) regardless of whether one is trying to facilitate the processes or trying to inhibit the process then it is assumed to be uncontrolled or automatic.

These methods allow examination, via the application of PDP methods pioneered by Jacoby (Jacoby, 1991; also see Govorun & Payne, 2006 for applications to social psychology), of the extent to which controlled processes and automatic processes might be involved in the generation of behavior predictions. In considering this idea, it is important to note that in Study 3 this procedure does *not* assess the extent to which initial inference-making is controlled and/or automatic: Instead, the procedure assesses the manner (controlled or automatic) in which these inferences, *once made*, affect the subsequent likelihood judgments.

From this perspective, one might not expect controlled processing to vary across instructional conditions. There is no obvious reason why the instructions that we provided ought to alter peoples' ability to intentionally control the use previously seen information to make behavior predictions. However, one might hypothesize that stronger initial inferences ought to produce stronger automatic influences on subsequent behavior predictions. Hence, if, relative to the other two instruction conditions, the lie detection instructions reduced initial inference-making, then it simply makes sense that automatic influences of initial inferences on subsequent behavior predictions also ought to decrease. In comparison, if subjects make spontaneous inferences in the familiarization condition in much the same way and to the same extent they do when explicitly instructed to make inferences, then the impact of automatic processing on subsequent behavior predictions should be equivalent in these two instructional conditions.

Method

Subjects

One hundred and thirty-nine subjects from Northern Illinois University participated. They volunteered for the study and received credit toward completion of a course research option as compensation. Data from twelve subjects were not used because there was no variance in their responses.[4]

Procedure

In the exposure task of Study 3, subjects were exposed to 16 photos paired with the TI behaviors (the TNI behaviors were not used in this study). Prior to viewing the actor-behavior pairs, some subjects received no impression formation instructions, and were simply told to *familiarize* themselves with the materials because they would be used later in the study. Other subjects received instructions *to try and form an impression* of the persons shown in the photos. Finally, other subjects were told to discern whether the behavior described by the actor in each photo was *the truth or a lie*.

After a short filler task, subjects were shown the 16 actor photos, each accompanied by new behaviors. As in Study 2, each photo appeared three times: Once with a new behavior that *matched the trait* implied by the behavior shown with the actor earlier, once with a new behavior that was *evaluatively consistent* with the implied trait, and once with a new behavior that was *evaluatively inconsistent* with the implied trait. Using the same 7-point scale described in Study 2, subjects rated the likelihood that the actor would perform each behavior.

[4] The data from the twelve omitted participants indicated a likelihood rating of '4' for all trials. These reaction times for these responses, recorded by the computer, were also identical for all responses for all twelve participants. Hence, we concluded that there was a computer error in recording responses for these subjects. This is what led us discard their data.

Eight novel photos were also shown during this stage; ratings were provided for new behaviors that could be performed by these actors. These behaviors were the same ones that accompanied the familiar actor photos. These ratings are thus a reflection of peoples' perceived likelihood of a behavior in the absence of a prior behavior and corresponding trait inference.

Prior to rating how likely an actor was to perform the new behaviors with which the actor was paired, half of the subjects were shown instructions (adapted from Uleman, Blader, & Todorov, 2005) that read:

> In the earlier task, pictures were paired with descriptions of behaviors that implied the actor had a certain personality trait. Some of those behavior descriptions were about the persons in the photos. For the next task, we will only use photos of people that were described by the behaviors seen earlier. The behaviors were provided by that person. Since the behaviors were about the person, they're informative when forming an impression of him or her. Try to remember the trait that was implied by the behaviors to help you in this next phase.

Essentially, subjects were instructed to use the trait information implied by the behavior shown with the actor earlier to inform their behavior predictions; therefore, following Jacoby (1991), this is referred to as the *inclusion task*.

After completing this task, these subjects read the following instructions:

> In the earlier task, pictures were paired with descriptions of behaviors that implied traits. Some of those behavior descriptions were NOT about the persons in the photos. They describe someone else and did not describe the person in the photo. We checked with the people to be sure. We paired them with photos randomly. Since the behaviors were NOT about the persons in the photos they are potentially misleading when trying to form an impression. Try to remember the traits implied by the behaviors so that you can ignore them and are not influenced by them.

In this task subjects were asked to *exclude* the potential influence of the behavior shown with the actor earlier from their predictions; therefore, following Jacoby (1991), we refer to this as the *exclusion task*. Ratings were made on the same response scale used in the inclusion task. After completing the final task, subjects were debriefed and dismissed.

To reduce confusion, only eight of sixteen actors shown during the exposure task were rated in each instruction condition. That is, of sixteen photos originally paired with trait-implicative behaviors, eight appeared in the inclusion task and the other eight appeared in the exclusion task. This was done to ensure that participants were not confused by trying to include information for an actor and then later trying to exclude information for the same actor.

Only half of the subjects completed the procedure in the order described in the paragraphs mentioned earlier; the other half of the subjects completed the exclusion task prior to the inclusion task. Analyses indicated that the order of instructions did not affect the results ($F<1$), so this manipulation was ignored in subsequent analyses.

Results and discussion

Predictions of future behavior were affected by new behavior type, pre-exposure task instructions, and task type

Average ratings for each new behavior that accompanied the novel control photos were calculated for each new behavior. These average baseline ratings were then subtracted from the rating for the same behavior that accompanied a photo of an actor who was previously described as performing a TI behavior. A positive difference score reflects a likelihood rating that is greater than the baseline mean for that specific new behavior.[5] These difference scores were entered into a 3 (Instruction: familiarization vs. impression goal vs. lie detection; between-subjects) × 3 (New Behavior Type: trait match vs. evaluatively consistent vs. evaluatively inconsistent; within-subjects) × 2 (Task Type: inclusion task vs. exclusion task: within-subjects) mixed-model ANOVA.

The analysis yielded three significant effects. As in Study 2, a new behavior type effect reflected both specificity and halo, $F(2,248) = 46.34$, $p<.0001$, $\eta^2 = .272$ (see Fig. 2). However, as revealed by the Instruction × New Behavior Type interaction, $F(4,248) = 5.24$, $p<.0001$, $\eta^2 = .078$, these effects were related to the mindset induced by the instructions provided prior to the exposure task. As the data in Fig. 2 show, subjects given lie detection instructions showed an especially small tendency to use the trait implications of a prior actor behavior to make predictions about future actor behavior. This reduction occurred most markedly for exact trait match behaviors, and occurred to a lesser extent for those behaviors that were only evaluatively consistent with, or evaluatively inconsistent with, the original behavior.

Another important element of this interaction concerned the impact of the impression formation instruction on behavior predictions. As can be noted from the data in Fig. 2, subjects given this instruction did not show evidence of enhanced trait inference-making relative to familiarization subjects. The similarity in the data yielded by these two conditions provide strong evidence that subjects in the familiarization condition were making spontaneous trait inferences about the actors from the actors' behaviors, and that these inferences (and the subsequent predictions) were essentially the same as those made by those who were explicitly told to form impressions. This becomes especially evident when we drop participants in the lie detection condition from the analyses. When only examining the data from those in the familiarization and impression formation condition the Instruction × New Behavior Type interaction becomes non-significant, $F(2,168) = .482$, *ns*. This not only shows that data from those in the lie detection condition were driving the significant interaction, it shows participants in the familiarization and impression formation conditions provided similar patterns of data.

The second significant effect of note that was yielded by the analysis was a Task Type × New Behavior Type interaction, $F(2,248) = 12.65$, $p<.0001$, $\eta^2 = .093$. The means for this interaction are presented in Fig. 3. They clearly show that subjects in the exclusion condition were able, at least to some extent, to use their ability to recall the previously seen behaviors and mute the extent to which those behaviors affected likelihood judgments. This effect is evident in both a reduction of the size of the difference scores reflecting the specificity effect in predictions, as well as in the magnitude of the halo effect. However, data also show that the effect of prior trait inference activity was not eliminated in the exclusion condition. This is especially evident in likelihood judgments for behaviors with trait implications that exactly matched the trait implications of original actor behaviors. This perception was verified via a *t* test that showed that the difference score for the trait-consistent behavior predictions was significantly greater than 0, $t(126) = 3.676$, $p<.0001$, even when subjects were trying to inhibit this information.

Assessing the extent to which controlled processing and automatic processing are involved in the influence of prior inferences on behavior predictions

We have argued that participants' behavior predictions can be affected by two qualitatively different sources of memory. First,

[5] Difference scores can be used to hide imperfections in the data. That is not the case here. We report results from difference score analyses because they lend simplicity to our report of results, and they provide a conceptual link to the process dissociation analyses that are reported later in this section.

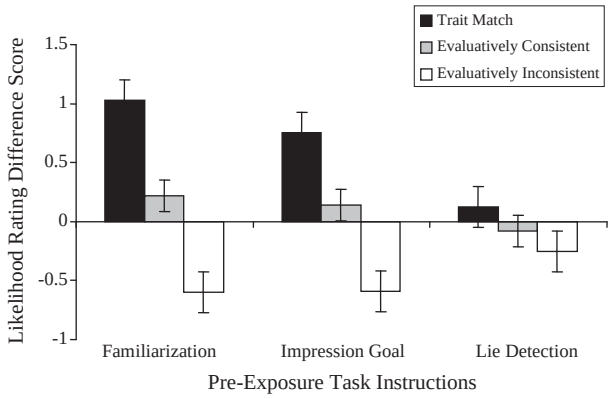


Fig. 2. *Note*. Mean new behavior likelihood ratings from Study 3; New Behavior Type × Instruction interaction.

participants can consciously recall and use information in a controlled manner to inform their behavior predictions by explicitly recalling the original behaviors or their original inferences. In Study 3, we had participants explicitly try to inhibit the potential influence of such information on their predictions. Subjects' inability to exclude the influence of the behaviors that they saw in the exposure task from their likelihood judgments suggests that some of this influence was outside of their control. Estimates of the extent to which automatic and controlled processes contributed to these predictions were more precisely obtained by applying Jacoby's (1991) process dissociation procedure to our data.

To do so, we first needed to convert the difference scores obtained for the trait-matching behaviors into *scaled difference scores*. This was accomplished by dividing each difference score by six, producing a score that reflects the magnitude of the difference score relative to the maximum difference score permitted by the scale (which is 7 − 1, or 6). For example, suppose that a subject provided a likelihood rating of 6 for an actor's new behavior *X*. If the baseline measure for *X* yielded a likelihood rating of 4, then the difference would be 2 (6 − 4 = 2). This difference score would then be divided by 6 to obtain the scaled difference score (2/6 = .333). In other words, a rating of 6 would represent 33.3% more of a "likelihood rating" than baseline.

For each subject, these scaled difference scores were separately calculated for the behavior predictions in the inclusion condition and in the exclusion condition. For the following analyses we assume that in the inclusion condition that likelihood ratings are driven by either intentionally recalling the original behaviors or inferences (a Controlled Process or C) or, when recall fails, by an automatic influence of STIs on predictions (an Automatic Process or A). Given these descriptions, predictions in the inclusion condition can be conceived of as reflecting C + A(1 − C). Conversely, predictions in the exclusion condition can be thought of in terms of the joint probability of an automatic process becoming expressed and the failure of inhibiting this influence, and can be mathematically expressed as A(1 − C).

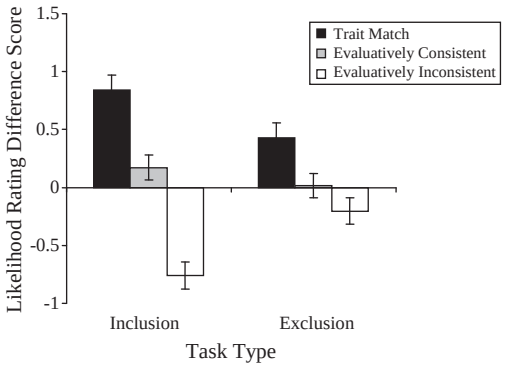


Fig. 3. *Note*. Mean new behavior likelihood ratings from Study 3; New Behavior Type × Task Type interaction.

To calculate the parameter estimate of controlled processes (C), the scaled difference scores obtained in the exclusion condition were subtracted from the scaled difference scores obtained in the inclusion condition (C = I − E; C scores below 0 were recorded to 0). Jacoby (1991) notes that the logic behind this calculation is that in the inclusion task subjects are trying to use the behavior information and in the exclusion task subjects are trying to inhibit the use of this information. To the extent that all this trying is successful, the scaled difference scores in the I condition should be greater than scaled difference scores in the E condition. Thus, if Jade kicking the cat led me to conclude that she was mean and I'm making predictions about the future likelihood of her yelling at her baby brother, recalling the original behavioral observation or inference should serve to increase the scaled difference score when successfully trying to use them and decrease the scaled difference score when successfully trying not to use them.

Conversely, to the extent that the original behavioral observation elicits a trait inference and this inference affects the likelihood judgments without participant awareness, then the scaled difference scores in the two conditions would be the same. If this were the case, then the effect of the trait inference on predictions of future behavior can be assumed to be out of volitional control. Adapting Jacoby's (1991) protocol, estimates of these automatic processes were mathematically derived. To do so, the scaled difference score in the exclusion condition was divided by the difference between 1 and the estimate of controlled processing (A = E/[1 − C]). Given this formula, one can see why it is necessary to recode negative estimates of the controlled processing parameter to the value of zero: Such a practice avoids artifactual reduction of the automatic processing parameter.

We collapsed the data across behaviors, calculating an average C score and A score for each subject. The calculated estimates of controlled processing for each subject were entered into a one-way ANOVA with instruction type as the independent variable. This variable failed to predict the C scores, $F(2,124) = 1.326$, *ns*. This result suggests that the controlled processes contributing to predictions of future behaviors were not greatly affected by the pre-exposure task instructions.

In comparison, results of the analysis of the automatic processing parameter estimates, $F(2,124) = 9.212$, $p < .01$, $\eta^2 = .129$, suggests that the extent to which inferences were automatically used to generate behavior predictions may have been affected by the instructions given

prior to the exposure task.[6] The means for this effect show that estimates of automatic processing were highest for those in the familiarization condition, ($M = .153$, $SD = .14$) middling for those receiving impression formation instructions ($M = .123$, $SD = .12$), and lowest for those receiving lie detection instructions ($M = .05$, $SD = .06$). Post-hoc comparisons suggested that this effect was largely driven by the low automatic processing parameter estimates given to those in the lie detection condition: They evinced lower automatic processing parameter estimates than those receiving familiarization instructions or those receiving impression formation instructions, both *Tukey's HSD* $p < .01$. The automatic processing did not differ between the familiarization and impression formation conditions, *Tukey's HSD* $p = .44$.

Some may argue that participants that had mean inclusion likelihood judgments below baseline showed no evidence of inference-making and the data cannot be dissected via a PDP analysis. If we restrict our analyses to only those with inclusion likelihood judgments greater than or equal to our baseline, we could exclude data from 38 participants from analyses ($n_{STI} = 8$, $n_{Impression\ Formation} = 12$, $n_{Lie\ Detection} = 18$). Two one-way ANOVAs were conducted on the data from the 89 remaining individuals. The implications of results of these analyses do not differ from the implications of the results reported earlier. Comparison of the controlled process parameter estimates do not differ among conditions, $F(2,86) = .279$, *ns*. However, the automatic process parameter estimates do differ among conditions, $F(2,86) = 6.938$, $p < .01$, $\eta^2 = .139$: The estimates in the familiarization ($M = .177$, $SD = .15$) and impression formation ($M = .137$, $SD = .12$) conditions do not differ, *Tukey's HSD* $p = .39$, but both are greater than the estimate in lie detection condition ($M = .054$, $SD = .07$), both *Tukey's HSD* $p < .01$.

In evaluating these results, it is important to recall that in this paradigm these parameter estimates do not refer to the extent to which initial trait inference-making is controlled and/or automatic: These estimates refer to the different sources of memory that could affect behavior predictions. One memory source is conscious recall for the behavior on which a prediction can be based. The other memory source is an implicit inference that exerts an influence outside of awareness. From this perspective, the fact that the estimates of C do not vary across the instructional conditions is not surprising. There is no obvious reason why the instructions that we provided ought to alter peoples' ability to intentionally control the use of previously seen behaviors. However, one might hypothesize that stronger initial inferences ought to produce stronger automatic influences on subsequent behavior predictions. Hence, if, relative to the other two inference conditions, the lie detection instructions reduced initial inference-making, then it simply makes sense that the estimate of the potency of automatic influence, embodied in the parameter estimate A, should also be lessened relative to the same parameter estimate observed in the other two instruction conditions. Using similar logic, if subjects make spontaneous inferences in the familiarization condition in much the same way and to the same extent that they do when explicitly instructed to make inferences, then one would expect the impact of automatic processing on subsequent behavior predictions to be equivalent in these two instructional conditions. Both of these outcomes emerged in Study 3.

General discussion

The automaker, Ford, featured golfer Phil Mickelson in an advertising campaign. In the campaign, predictions about Phil's future behaviors (e.g., Mickelson singing Toby Keith's song "Who's Your Daddy?") depended on a prior trait (adventurous) that had been abstracted from observing Phil's prior behaviors. The advertisements included a voice-over indicating that: "Phil Mickelson will try just about anything…which leads to the question: What will Phil do next?" Each ad then humorously depicted Mickelson's adventurous nature and willingness to try new things.

Research results confirm that, as exemplified by the Mickelson ads, people use their prior inferences to make predictions of future behavior (Dunning et al., 1990; Griffin et al., 1990; Newman, 1991, 1996; Nussbaum et al., 2003). The studies reported in Studies 1 and 2 extend these prior results by showing that these predictions occur even when traits are spontaneously inferred. The results from Studies 2 and 3 suggest that the basis of the behavior predictions were trait inferences, not evaluative inferences. Moreover, in Study 3 the extremity of the behavior predictions made from spontaneous inferences did not differ from extremity of the inferences made after inferences were intentionally made, a result that suggests that the trait inferences made in the two conditions were similar or identical. Finally, the results obtained in Study 3 show that these behavior prediction effects may partially be a consequence of the strength of the initial inference and can occur even when people may be trying to stop them.

Such effects are of considerable potential importance. One theoretical implication is that there may be limits on the extent to which one can effectively apply correction processes (see Wegener, Petty, & Dunn, 1998) in the context of prior inferences when considering an individual's potential future actions. Correction processes typically require a perceiver to be aware of the influence of the knowledge that they possess and also be motivated to correct for its influence: The current studies suggest people may not be aware of the full impact of their prior inferences, and therefore, might not be able to fully correct for their influence. As the hall-of-fame pitcher Walter Johnson quipped, "You can't hit what you can't see." Thus, observers who form an impression of an actor may be unable to disregard that impression when making predictions of future behavior, even when trying to do so—even when they know that the initial impression was erroneous.

Such effects speak to the power of trait inference-making, and make it even more important to understand when such inferences are made and the properties that such inferences possess. The data from our studies add to the corpus of evidence suggesting that people readily make trait inferences about others, and need no prompting to do so. These data also suggest that spontaneous inferences behave in ways resembling the intentional judgments studied by researchers examining intentional inferences, as occurred in explorations of Correspondent Inference Theory (e.g., Jones & Davis, 1965). One of the characteristics observed is that trait inferences act as if they have causal force, in that they are used to construct predictions of future behavior. The fact that spontaneous inferences have such properties makes it difficult to claim that the results from earlier STI research do not reflect inferences, but merely reflect the action of labels that are associated to constructs (e.g., Bassili, 1989; Brown & Bassili, 2002).

However, some might claim that our conclusion that spontaneously inferred traits possess causal properties, although plausible, may not be warranted by the current data. They might argue that if trait *X* merely becomes associated with an actor in memory, the trait may serve as a cue as to how the actor may behave in the future. In this regard, it would be especially interesting to see if the same predictive effects that were observed in this study also occur for the phenomenon of spontaneous trait transference (STT; Mae, Carlston, & Skowronski, 1999; Skowronski, Carlston, Mae, & Crawford, 1998). This phenomenon shows that an informant who describes a trait-implicative behavior performed by another person becomes associated with that trait. We suspect that the effects of this association on predictions of the informant's behavior may be limited. Such a finding would add to the corpus of findings suggesting that trait inferences about an individual go beyond mere associations between that individual and trait knowledge (Crawford, Skowronski & Stiff, 2007; Crawford, Skowronski, Stiff, & Leonards, 2008; Crawford et al., 2007).

[6] The automatic estimates and the controlled estimates cannot be used as a within-subjects variable in an ANOVA because one is calculated from the other; this dependency violates the assumptions of the ANOVA.

However, we would also argue that such data may not be necessary to combat the association-only claim. One of the consistent findings from the studies that compare inferences and associations in the context of the STT vs. STI studies is that the associative effects observed in STT conditions have so far *always* been considerably weaker than the inferential effects observed in the STI conditions (see Skowronski et al., 2008). This fact makes it difficult to accept an argument that the predictions made in Study 3's familiarization condition were association-based. The fact that the strength of such predictions did not differ from the strength of the predictions made by those who explicitly formed trait impressions of the actors simply does not fit with the usual inference vs. association pattern that has consistently characterized this line of research.

A similar conclusion can be derived from Study 3's halo effect data. Carlston and Skowronski (2005) argue that the emergence of a halo effect is one of the hallmarks that seem to distinguish the formation of inferential links between actors and traits from the formation of links that are purely associative. This assertion is backed by the fact that halo effects frequently occur when spontaneous inferences are made; they almost never are seen when individuals are merely associated with traits. Tellingly, in Study 3 halo effects emerged, and the halo effects observed in the familiarization condition were not significantly different from the halo effects observed when people intentionally formed trait inferences. This again points to equivalence in the processes and structures that were driving the results in these two conditions.

A final point to be drawn from the studies described in this article is that the downstream consequences of trait inferences can operate independently from recall for the information on which such inferences are based: Exposure to trait-implicative behaviors influenced prediction of future actor behavior, even when subjects had little memory for the original trait-implicative behaviors.

However, some may argue that there are reasons to be cautious about this conclusion. First, people may not naturally encode their social worlds in a manner that is amenable to verbatim recall; therefore, asking subjects to recall previously encountered behavior descriptions verbatim may not best capture the status of their memory for the behaviors. Second, subjects attempted to recall the behaviors at the end of the experiment session, and the timing of this recall session may not adequately capture the status of their behavior memories at the time they made their behavior predictions. Hence, some may argue that because our memory measure may not have adequately captured the content of peoples' memory for the original behaviors, the results of our analyses may minimize the relation between memory for the original behaviors and behavior predictions. We agree that our free recall measure is very conservative with respect to the assessment of the contents of memory. However, in this regard, we also note that at least one study has used a recognition task to probe behavior memories in these kinds of tasks (see Carlston & Skowronski, 1994). The conclusion from the data collected for that study was the same as the conclusion derived from the studies that have used the photo-cued recall task: The memory-mediated mechanism was implausible.

Moreover, we note that the important question is not the extent to which participants retain memory traces of the original behaviors that they viewed, but it is the extent to which they can intentionally access and use those memories to generate behavior predictions. Given this question, one can argue that that the photo-cued recall task may actually be the *most appropriate measure* of memory to use when evaluating the notion that behavior predictions were mediated by recall of the behavior that was originally paired with an actor photo. The reason is simple: For the mediated process to work, when cued with an actor photo a subject would need to first recall the behavior when cued by the behavior. The mediated process cannot explain the behavior predictions that we obtained if people cannot substantially recall the behavior, or some reasonable facsimile of the behavior, when cued by the actor photo. Given the poor photo-cued recall performance of our subjects, the behavior recall-mediated route to behavior predictions simply seems quite implausible.

Coda

People live in social environments. Effective navigation of these environments may characterize those who make fast and usable inferences about others. Indeed, data show that even a cursory interaction may sometimes provide enough information to form accurate impressions of others (e.g., Ambady, Bernieri & Richeson, 2000). Certainly, we would not argue that peoples' trait inferences and subsequent behavioral predictions are always accurate—to do so would ignore the research showing that inferences and judgments are sometimes biased (e.g., Dunning et al., 1990; Gilbert, 1989; Jones & Harris, 1967; Trope, 1986). However, spontaneous inferences and the predictions on which they depend need not always be accurate; they need only to be accurate often enough that an advantage is given to those who sometimes make accurate inferences and predictions.

The advantages of accurate and rapid inference-making can be couched in alternative theoretical frames. In evolutionary theory, the advantages could lead to reproductive success. In economic terms, the advantages would reduce costs and increase benefits. Regardless, given the important need for rapid and accurate person perception, it may not be especially surprising that humans have become proficient at seeing meaning in others' behaviors. If accurate, these impressions can serve to effectively direct subsequent interactions with those others, in part by allowing accurate predictions of others' behaviors.

However, many questions remain about spontaneous inferences and the extent to which they affect behavior predictions. For which perceivers do they occur most strongly? Do they occur for some kinds of social targets more than for other kinds of targets? For which behaviors? For which traits? Does it matter whether the inference is generated passively or is generated in an attempt to explain a behavior (as might be suggested by Hamilton, 1988; 1998)? While extant research has made progress in addressing these kinds of questions (e.g., Uleman, Winborne, Winter & Schechter, 1986; Wigboldus, Dijksterhuis, & van Knippenberg, 2003) much remains to be done in this domain. In the near future, we hope to be counted among those researchers who will be reporting results relevant to these important questions.

Acknowledgments

We thank Brad Sagarin, Amanda Durik, Keith Payne, and several anonymous commentators for their helpful comments. We especially thank Dave Hamilton for his willingness to intensively discuss the theoretical implications of the studies described in this research. We also thank the undergraduate research assistants who reside in the NIU Social Cognition Lab for their data collection and data coding efforts.

References

Ambady, N., Bernieri, F. J., & Richeson, J. A. (2000). Toward a histology of social behavior: Judgment accuracy from thin slices of the behavioral stream. *Advances in Experimental Social Psychology*, *32*, 201–271, doi:10.1016/S0065-2601(00)80006-4.

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, N.J.: Erlbaum.

Biner, P. M., Angle, S. T., Park, J. H., Mellinger, A. E., & Barber, B. C. (1995). Need state and the illusion of control. *Personality and Social Psychology Bulletin*, *21*, 899–907, doi:10.1177/0146167295219004.

Brown, R. D., & Bassili, J. N. (2002). Spontaneous trait associations and the case of the superstitious banana. *Journal of Experimental Social Psychology*, *38*, 87–92, doi:10.1006/jesp. 2001.1486.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, *64*, 123–152.

Carlston, D. E., & Skowronski, J. J. (1994). Savings in the relearning of trait information as evidence for spontaneous inference generation. *Journal of Personality and Social Psychology*, *66*, 840–856, doi:10.1037/0022-3514.66.5.840.

Carlston, D. E., & Skowronski, J. J. (2005). Linking versus thinking: Evidence for the different associative and attributional bases of spontaneous trait transference and spontaneous

trait inference. *Journal of Personality and Social Psychology*, *89*, 884–898, doi:10.1037/0022-3514.89.6.884.
Chakrabarti, B., & Baron-Cohen, S. (2008). The biology of mind reading. In N. Ambady, & J. J. Skowronski (Eds.), *First impressions* (pp. 57–86). New York: Guilford.
Crawford, M. T., Skowronski, J. J., & Stiff, C. (2007a). Limiting the spread of spontaneous trait transference. *Journal of Experimental Social Psychology*, *43*, 466–472, doi:10.1016/j.jesp. 2006.04.003.
Crawford, M. T., Skowronski, J. J., Stiff, C., & Leonards, U. (2008). Seeing but not thinking: Limiting spontaneous trait transference II. *Journal of Experimental Social Psychology*, *44*, 840–847, doi:10.1016/j.jesp. 2007.08.001.
Crawford, M. T., Skowronski, J. J., Stiff, C., & Sherer, C. R. (2007b). Interfering with inferential, but not associative, processes underlying spontaneous trait inference. *Personality and Social Psychology Bulletin*, *33*, 677–690, doi:10.1177/0146167206298567.
Dunning, D., Griffin, D. W., Milojkovic, J., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, *58*, 568–581, doi:10.1037/0022-3514.58.4.568.
Empirisoft (2004). DirectRT (Version 2004.1.0.42) [Computer software]. New York, NY: Empirisoft.
Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. *Cognition and Emotion*, *15*, 115–141, doi:10.1080/0269993004200024.
Gilbert, D. T. (1989). Thinking lightly about others: Automatic components of the social inference process. In J. Uleman, & J. A. Bargh (Eds.), *Unintended thought* (pp. 189–211). New York: Guilford.
Govorun, O., & Payne, B. K. (2006). Ego depletion and prejudice: Separating automatic and controlled components. *Social Cognition*, *24*, 111–136, doi:10.1521/soco.2006.24.2.111.
Griffin, D. W., Dunning, D., & Ross, L. (1990). The role of construal processes in overconfident predictions about self and others. *Journal of Personality and Social Psychology*, *59*, 1128–1139, doi:10.1037/0022-3514.59.6.1128.
Hamilton, D. L. (1988). Causal attribution viewed from an information-processing perspective. In D. Bar-Tal, & A. W. Kruglanski (Eds.), *The social psychology of knowledge* (pp. 359–385). Cambridge, UK: Cambridge University Press.
Hamilton, D. L. (1998). Dispositional and attributional inferences in person perception. In J. M. Darley, & J. Cooper (Eds.), *Attribution and social interaction: The legacy of Edward E. Jones* (pp. 99–114). Washington, D.C.: American Psychological Association.
Hassin, R. R., Bargh, J. A., & Uleman, J. S. (2002). Spontaneous causal inferences. *Journal of Experimental Social Psychology*, *38*, 515–522, doi:10.1016/S0022-1031(02)00016-1.
Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541, doi:10.1016/0749-596X(91)90025-F.
Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 220–266). New York: Academic Press.
Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, *3*, 1–24.
Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 192–241). Lincoln: University of Nebraska Press.
Klein, S., Cosmides, L., Murray, E., & Tooby, J. (2004). On the acquisition of knowledge about personality traits: Does learning about the self engage different mechanisms than learning about others? *Social Cognition*, *22*, 367–390, doi:10.1521/soco.22.4.367.38295.
Komorita, S. S., & Parks, C. D. (1995). Interpersonal relations: Mixed-motive interaction. *Annual Review of Psychology*, *46*, 183–207, doi:10.1146/annurev.ps.46.020195.001151.
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, doi:10.1037/0022-3514.77.2.233.
Nesse, R. M. (2005). Natural selection and the regulation of defenses: A signal detection analysis of the smoke detector principle. *Evolution and Human Behavior*, *26*, 88–105.
Newman, L. S. (1991). Why are traits inferred spontaneously? A developmental approach. *Social Cognition*, *9*, 249–269.
Newman, L. S. (1996). Trait impressions as heuristics for predicting future behavior. *Personality and Social Psychology Bulletin*, *22*, 395–411, doi:10.1177/0146167296224006.
Nussbaum, S., Trope, Y., & Liberman, N. (2003). Creeping dispositionalism: The temporal dynamics of behavior. *Journal of Personality and Social Psychology*, *84*, 485–497, doi:10.1037/0022-3514.84.3.485.
Pittman, T. S., & Heller, J. F. (1987). Social motivation. *Annual Review of Psychology*, *38*, 461–489, doi:10.1146/annurev.ps.38.020187.002333.
Schaller, M. (2008). Evolutionary bases of first impressions. In N. Ambady, & J. J. Skowronski (Eds.), *First impressions* (pp. 15–34). New York: Guilford Press.
Schaller, M., Park, J. H., & Kenrick, D. T. (2007). Human evolution and social cognition. In R. I. M. Dunbar, & L. Barrett (Eds.), *Oxford handbook of evolutionary psychology* (pp. 491–504). Oxford, UK: Oxford University Press.
Schaller, M., Simpson, J. A., & Kenrick, D. T. (2006). *Evolution and social psychology*. New York: Psychology Press.
Shoda, Y., & Mischel, W. (1993). Cognitive social approach to dispositional inferences: What if perceiver is a cognitive social theorist? *Personality and Social Psychology Bulletin*, *19*, 574–585, doi:10.1177/0146167293195009.
Skowronski, J. J., Carlston, D. E., & Hartnett, J. L. (2008). Spontaneous impressions derived from observations of behavior: What a long, strange trip its been (and it's not over yet). In N. Ambady, & J. J. Skowronski (Eds.), *First impressions* (pp. 313–333). New York, NY: Guilford Press.
Skowronski, J. J., Carlston, D. E., Mae, L., & Crawford, M. T. (1998). Spontaneous trait transference: Communicators take on the qualities they describe in others. *Journal of Personality and Social Psychology*, *74*, 837–848, doi:10.1037/0022-3514.74.4.837.
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*, *39*, 549–562, doi:10.1037/0022-3514.83.5.1051.
Trope, Y. (1986). Identification and inference processes in dispositional attributions. *Psychological Review*, *93*, 239–257.
Uleman, J. S., Blader, S. L., & Todorov, A. (2005). Implicit impressions. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 362–392). New York: Oxford University Press.
Uleman, J. S., Newman, L. S., & Moskowitz, G. B. (1996). People as flexible interpreters: Evidence and issues from spontaneous trait inferences. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 211–279). San Diego, CA: Academic Press.
Uleman, J. S., Saribay, S. A., & Gonzalez, C. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual Review of Psychology*, *59*, 329–360, doi:10.1146/annurev.psych.59.103006.093707.
Uleman, J. S., Winborne, W. C., Winter, L., & Schechter, D. (1986). Personality differences in spontaneous personality inferences at encoding. *Journal of Personality and Social Psychology*, *51*, 396–403, doi:10.1037/0022-3514.51.2.396.
Wegener, D. T., Petty, R. E., & Dunn, M. (1998). The metacognition of bias correction: Naive theories of bias and the flexible correction model. In V. Yzerbyt, G. Lories, & B. Dardenne (Eds.), *Metacognition: Cognitive and social dimensions* (pp. 202–227). London: Sage.
Wigboldus, D. H., Dijksterhuis, A., & van Knippenberg, A. (2003). When stereotypes get in the way: Stereotypes obstruct stereotype-inconsistent trait inferences. *Journal of Personality and Social Psychology*, *84*, 470–484, doi:10.1037/0022-3514.84.3.470.
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, doi:10.1037/0022-3514.47.2.237.
Zajonc, R. B. (1984). On the primacy of affect. *The American Psychologist*, *39*, 117–123, doi:10.1037/0003-066X.39.2.117.