Journal of Experimental Social Psychology 47 (2011) 37–49

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

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

The automatic and co-occurring activation of multiple social inferences

Andrew R. Todd [a,*], Daniel C. Molden [b], Jaap Ham [c], Roos Vonk [d]

[a] University of Cologne, Germany
[b] Northwestern University, USA
[c] Eindhoven University of Technology, The Netherlands
[d] Radboud University Nijmegen, The Netherlands

ARTICLE INFO

Article history:
Received 18 September 2009
Revised 25 July 2010
Available online 14 August 2010

Keywords:
Automaticity
Spontaneous inference
Person perception
Attribution

ABSTRACT

People make a variety of automatic inferences when observing others' actions. These include inferences about stable dispositions as well as transitory goal states and social situations. However, models of social inference have rarely considered whether different types of automatic inferences can co-occur. We present three experiments in which participants were incidentally exposed to texts depicting behaviors that afforded inferences about actors' traits and the social situations these actors were experiencing. Results from lexical decision and probe-recognition tasks revealed heightened activation of both trait and situational inferences; furthermore, this co-occurring activation was spontaneous, unconscious, and independent of processing resources or specific impression-formation goals. A fourth experiment extended these findings by showing that when participants were asked to make deliberate attributional judgments of the same set of behaviors, typical goal-directed biases reflecting the selection of either trait or situational interpretations emerged. Implications for social inference processes are discussed.

Imagine happening upon a colleague struggling mightily to operate a new cell phone. Hiding a smirk, you ponder the depths of his technical ineptitude with a device that most children can easily master. However uncharitable it may be, this impression would likely spring to mind just as readily as it occurs to you that you are wet when caught in the rain. In other words, people typically experience social perception as a "given"—as immediate and valid as their perception of their physical environments (Heider, 1958).

A long history of social perception research has revealed that our impressions of others emerge from a variety of distinct psychological processes, many of which occur quite automatically. For instance, numerous studies have shown that social perceivers automatically draw inferences about personality traits from brief observations of behavior (for a review, see Uleman, Saribay, & Gonzalez, 2008). More recently, studies have shown that other types of inferences, such as the goals an actor is pursuing (Hassin, Aarts, & Ferguson, 2005) or the current pressures of the situation an actor is experiencing (Ham & Vonk, 2003; Lupfer, Clark, & Hutcherson, 1990), can arise automatically as well. Thus, upon observing the frustrations of your colleague, you might also infer that he intends to send or retrieve an urgent message or that the cell phone is not working properly.

The seeming ease with which people generate inferences about others' traits, goals, and social situations raises several important questions: first, although perceivers appear capable of making a variety of different automatic social inferences, do these automatic inferences ever co-occur, or is a single inference activated among the different possibilities? Second, if multiple automatic inferences are activated concurrently, then what accounts for the prevalence of robust attributional biases – based on people's current inference goals, lay theories, and available cognitive resources – that are frequently observed in social perception (Gilbert & Malone, 1995; Krull & Erickson, 1995a; Molden & Dweck, 2006)? The current research was designed to investigate these questions.

The automaticity of social inferences

Research on automaticity has defined the criteria for considering particular psychological processes to be "automatic" in a variety of ways (see Moors & De Houwer, 2006). However, the term automatic is most often used to describe any process that is characterized by one or more of a limited set of conceptually separable features: it is activated *spontaneously* (without intention), occurs *unconsciously* (outside of awareness), is *efficient* (requires few cognitive resources), or is *uncontrollable* (endures in the face of efforts to inhibit it; Bargh, 1994).

Abundant research has shown that merely attending to behaviors, in the absence of an intentional impression-formation goal, can elicit spontaneous trait inferences (e.g., Uleman & Moskowitz, 1994; see Uleman et al., 2008). When directly asked, people report little to no awareness of having made such inferences (Winter & Uleman, 1984; Winter, Uleman, & Cunniff, 1985), yet they continue to show evidence of having done so, even after explicit memory for the behaviors has decayed (Carlston & Skowronski, 1994; Carlston, Skowronski, & Sparks, 1995). Furthermore, these unintentional and unconscious trait inferences persist when people's cognitive resources are depleted (Todorov

* Corresponding author. Universität zu Köln, Department Psychologie, Sozialpsychologie I, Richard-Strauss-Str. 2, 50931 Köln, Germany. Fax: +49 221 470 1216 0.
E-mail address: atodd@uni-koeln.de (A.R. Todd).

0022-1031/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jesp.2010.08.006

& Uleman, 2003; Winter et al., 1985), even, it seems, when people are explicitly instructed to disregard the trait implications of the behavior (Uleman & Blader, 2002, cited in Todorov & Uleman, 2003; but see Uleman, Hon, Roman, & Moskowitz, 1996). Thus, trait inferences appear to occur with a high degree of automaticity.

The automaticity of trait inferences has long bolstered theoretical explanations for the *correspondence bias*, the tendency for perceivers to provide dispositional attributions for behaviors even when those behaviors are constrained by situational factors (Gilbert & Malone, 1995). Yet, other social inferences have also been found to exhibit some aspects of automaticity. For example, merely attending to descriptions of behavior can also elicit the spontaneous activation of goal-related concepts (Hassin et al., 2005) and of concepts describing the actor's situation (Ham & Vonk, 2003; Lupfer et al., 1990). Furthermore, in one line of research, Ham and Vonk (2003) demonstrated that the extent to which inferences about social situations were spontaneously activated by particular descriptions of behavior did not differ from the extent to which the relevant trait inferences were also spontaneously activated by the same descriptions.

Automatic and co-occurring social inferences

Findings suggesting that automatic social inferences are not necessarily limited to personality traits potentially undermine assumptions that some theorists have made regarding the privileged status of trait inferences at the earliest stages of attributional judgment (Gilbert & Malone, 1995; Ross & Nisbett, 1991). However, before these assumptions can truly be questioned, more conclusive evidence is required showing that other types of social inferences possess a similar level of automaticity as trait inferences. It is also necessary to empirically demonstrate that these other social inferences can automatically *co-occur* with trait inferences. That is, even if situational and goal inferences are activated automatically, they still could be overshadowed or inhibited in instances when trait inferences are also activated, which would reaffirm the primacy of traits in attributional judgment.

The idea that descriptions of behavior automatically elicit multiple, co-occurring inferences is generally consistent with research on text comprehension (e.g., Graesser, Lang, & Roberts, 1991; Kintsch, 1988) and with Read and Miller's (1993, 1998) neural network model of social perception (see also Reeder, 2009; Trope, 1986). Moreover, Ham and Vonk (2003) have provided preliminary evidence consistent with the automatic and co-occurring activation of trait and situational inferences. However, this evidence relied on comparing spontaneous trait inferences and spontaneous situational inferences between separate groups of participants. Consequently, it has yet to be conclusively demonstrated that individuals themselves automatically activate trait and situational inferences concurrently. Furthermore, whereas Ham and Vonk (2003) primarily examined the spontaneity with which trait and situational inferences occurred, they did not assess how unconscious or efficient these inferences were. Thus, our primary objective was to document the automatic co-occurrence of trait and situational inferences and to investigate the implications of this automatic co-occurrence on the social inference process more generally.

Although there are numerous inferences that social perceivers presumably activate when observing behavior (e.g., goals, beliefs, intentions; Malle, 1999; Read & Miller, 1993,1998; Reeder, 2009), we investigated co-occurring trait and situational inferences for two reasons: (a) the contrast between trait and situational inferences has been a primary focus in previous social inference research, and (b) by and large, trait and situational inferences are thought to be incompatible with each other (Gilbert & Malone, 1995; Heider, 1958; Kelley, 1973). Examining the automatic co-occurrence of trait and situational inferences thus affords direct comparisons to existing theories of social inference processes and provides a strong test case for demonstrating the existence of such co-occurrence. Before describing this investigation in detail, we first consider the implications that the automatic co-occurrence of trait and situational inferences would have for current accounts of social inference processes.

Automatic social inferences and biases in attributional judgment

Traditional accounts of social inference have identified several different processes that contribute to attributional judgments, each of which can engender biased impression formation (Gilbert & Malone, 1995; Krull & Erickson, 1995a; Lieberman, Gaunt, Gilbert, & Trope, 2002; Trope, 1986). Although these accounts differ in both the number and the operation of processes involved in social inference, they generally maintain that upon observing someone's actions, perceivers (a) decide what kind of action it is and form an initial interpretation of what this action means, and then (b) consider whether any other interpretations might also apply. Furthermore, whereas the process of forming an initial interpretation is conceptualized as requiring minimal cognitive resources, the process of considering alternative interpretations is thought to be more effortful.

According to this general perspective on social inference, biases typically arise from two primary sources. The first source of bias involves influences on perceivers' initial interpretations of behavior. For example, when actors' personality traits are a salient interpretation for their behaviors (which, based on the evidence for automatic trait inferences, has been assumed to be the typical case) or when perceivers deliberately pursue goals to form trait impressions, trait-based inferences often dominate their attributional judgments. In contrast, when the social situation an actor is facing is particularly salient or when perceivers deliberately pursue goals to form situational inferences, situation-based attributional judgments usually prevail (Gilbert, Pelham, & Krull, 1988; Krull, 1993; Krull & Erickson, 1995b; Lupfer et al., 1990; Molden, Plaks, & Dweck, 2006; Trope & Gaunt, 2000). However, the effects of differential initial interpretations on final attributional judgments also depend upon a second source of bias: the extent to which perceivers possess the cognitive resources needed to thoroughly consider alternative interpretations. The biasing effect of salient interpretations and specific inference goals is strongest when perceivers' cognitive resources are limited and they are less willing or able to integrate alternatives (Gilbert et al., 1988; Krull, 1993; Krull & Erickson, 1995b; Molden et al., 2006; Trope & Gaunt, 2000).

On the whole, most current accounts of social inference processes thus explicitly assume that considering multiple interpretations of behavior concurrently is a resource-dependent process that typically occurs only when perceivers are willing and able to dedicate resources to the task (but see Read & Miller, 1993, 1998). These accounts further assume that the biases in social inference that emerge when cognitive resources are scarce arise from the constrained activation or realization of different interpretations of behavior. Evidence for the automatic and co-occurring activation of multiple social inferences would question both of these assumptions and invite new analysis of the processes responsible for the robust biases that have been observed.[1]

[1] Although the social inference models proposed by Trope and colleagues (for a review, see Trope & Gaunt, 2003) do not explicitly discuss the automatic and co-occurring activation of multiple inferences, there are features of these integration models that are compatible with our perspective. For example, such models also include an earlier stage during which multiple pieces of dispositional and situational information might initially be spontaneously activated by, and thus influence, the interpretation of actors' behaviors, which is analogous to our position that perceivers consider multiple interpretations of behavior even when their cognitive resources are constrained. However, in contrast to our proposal that this activation represents the co-occurrence of multiple, yet distinct, automatic inferences about a particular behavior, Trope and colleagues primarily characterize this initial activation – identification in their terms – as a process whereby salient information about an actor's situation or prior disposition can automatically affect how this information is integrated into a single dispositional inference (Trope, 1986; Trope & Alfieri, 1997; Trope & Gaunt, 2000). That is, whereas Trope and colleagues' research has explored how salient information about an actor's situation can automatically affect inferences about his or her traits by disambiguating what that actor's behavior signifies, the current research focuses on whether multiple distinct inferences about an actor's traits and situation are automatically activated from descriptions of behavior.

The findings by Ham and Vonk (2003) suggesting that trait and situational inferences can spontaneously co-occur led them to speculate that the automatic activation of multiple social inferences represents an even earlier stage in the inference process—one that occurs before a definite initial interpretation of behavior is made and that is independent of perceivers' particular inference goals. To account for the well-documented biases that appear in people's resource-independent interpretations, Ham and Vonk further proposed an intermediary stage wherein perceivers make a selection among the active inferences. This distinction between activation and selection resembles a similar distinction made by Read and Miller (1993, 1998), who argue that perceivers, upon observing a behavior, initially activate a number of potential explanations, after which principles of explanatory coherence (Thagard, 1989) and augmenting and discounting (Kelley, 1973) guide the selection among the possible alternatives (see also Carlston & Skowronski, 2005). However, because few empirical investigations have attempted to separate these activation and selection processes, a secondary objective of the current research was to provide preliminary evidence for the operation of a selection process that follows the automatic activation of multiple inferences.

Overview of the current research

Experiments 1–3 were designed to provide evidence for the co-occurring activation of both trait and situational inferences that are spontaneous, unconscious, efficient, and independent of people's inference goals. In these experiments, participants read behavior descriptions that afforded both trait and situational inferences. In Experiment 1, participants had the goal of merely attending to the descriptions, and we assessed accessibility for trait-related and situation-related concepts using a lexical decision task. In Experiment 2, participants had the goal of memorizing the behavior descriptions, and we assessed the spontaneity and efficiency of trait and situational inferences using a probe-recognition task while some participants experienced a cognitive load. In Experiment 3, participants had the explicit goal of judging either actors' traits or their situations, and the spontaneity and efficiency with which they continued to form both types of inference were assessed in the same manner as Experiment 2.

Experiment 4 was designed to explore the relationship between the automatic and co-occurring activation of multiple social inferences observed in Experiments 1–3 and the biased attributional judgments that often subsequently arise. Using the same behavior descriptions, inference goal instructions, and cognitive load manipulation as Experiment 3, Experiment 4 examined whether *deliberate judgments* of actors' traits and situations do indeed diverge from the *incidental activation* of trait and situational inferences observed under identical conditions.

Experiment 1: Spontaneous and unconscious co-occurring inferences

The primary objective of Experiment 1 was to provide an initial test of whether perceivers automatically activate multiple, co-occurring social inferences. Participants read briefly presented descriptions of behavior, after which they completed a lexical decision task that included a word corresponding to a personality trait implied by the description, a word corresponding to a situation evoked by the description, and matched-control words. These procedures allowed us to assess the extent to which participants were faster to respond to both the trait words and the situation words in comparison to the matched-control words, a pattern indicative of co-occurring, automatic inferences.

Method

Participants

Participants were 40 native English-speaking students (25 women, 15 men) at a private university in the United States. There were no gender effects in any of the experiments in this article.

Procedure

As part of a study on "linguistic processing," participants completed a lexical decision task, during which they were asked to determine as quickly and accurately as possible whether or not a series of letter strings comprised actual English words. As a secondary "distracter task," participants read a single sentence and rated how interesting the sentence was (Hassin et al., 2005) before each block of letter-string judgments. Each block of lexical decision trials began with a sentence presentation (3000 ms), followed by a prompt to indicate how interesting the sentence was (1 = *not at all*, 7 = *extremely*). The rating task merely served to bolster the cover story; these ratings were not recorded and thus were not analyzed. A row of Xs then appeared for 500 ms (to signal the onset of a set of lexical decision trials) and was replaced by a letter string, which remained on the screen until participants classified it as a word or a non-word by pressing one of two keys labeled "Yes" and "No," respectively. Once participants had responded, the next block of trials began. Following the final block of trials, participants were questioned about any explicit attempts (a) to determine the causes of the behaviors, (b) to form an impression of the actor, or (c) to form an impression of the situation facing the actor.

Stimuli

In reality, the "distracter" sentences consisted of behavior descriptions that had been pre-tested to afford both trait and situational inferences. A basic requirement for assessing co-occurring inferences is a set of stimuli for which multiple inferences are equally available; before testing whether multiple inferences are automatically activated, it was first necessary to confirm that these inferences could reasonably be applied to the behavior (Higgins, 1996). In an initial pilot study, a group of 34 participants read 20 behavior descriptions selected from those used by Ham and Vonk (2003) and Uleman (1988) and listed two possible inferences about the traits of the actor and two possible inferences about the situation or context the actor was experiencing. From this initial set, we selected for the main study 8 sentences that elicited strong agreement (i.e., among at least 50% of participants) for both the trait and situational inferences. Thus, all behavior stimuli were (a) clear and unambiguous in the inferences they afforded and (b) interpretable in multiple ways.

The most frequently listed trait and situational inferences from the pilot study served as the critical stimuli in the lexical decision task. Each of these 8 trait and 8 situation words was matched with a control word that was equal in length and lexical frequency. We also created a set of pronounceable non-words (matched in length) to ensure an equal number of word and non-word trials within the task. To further ensure that the trait and situation words were properly matched with the control words, we conducted a second pilot study, during which a new group of 27 participants completed an identical lexical decision task without having read the behavior descriptions beforehand. Results revealed no a priori differences in response latencies based on word type (trait vs. situation vs. matched-control; $F<1$); separate comparisons revealed no differences between trait and control words or between situation and control words ($ps>.20$).

Results

Data from seven participants who reported explicitly trying to infer the causes of the behaviors were excluded, although retaining their data did not change the significance of any of the results. The remaining 33 participants did not report any intention or awareness of having made inferences; thus, their inferences can be surmised to be spontaneous and unconscious.

Participants' lexical decision latencies served as our primary index of automatic inference. Shorter latencies for both trait and situation words as compared to matched-control words would indicate that both types of inferences had been spontaneously activated while reading the behavior descriptions (Hassin et al., 2005). In addition, if the activation of these inferences is indeed co-occurring, then the latency advantage for both trait and situation words should remain equivalent across (a) the time that passes between exposure to the descriptions and inference assessment (i.e., whether the probe words appeared on trials 1–8 in the lexical decision task), and (b) the specific order in which trait and situation responses are assessed relative to each other. That is, if trait and situational inferences are truly co-active, then we would not expect to find that (a) the activation of one of these inferences appears or disappears more quickly than the other (cf. Swinney, 1979), or (b) that assessing one type of inference first (and perhaps making this inference particularly salient) somehow inhibits the other. Both scenarios would imply the presence of some antagonistic or inhibitory mechanism (Gilbert & Malone, 1995; Lieberman et al., 2002; Read & Miller, 1993, 1998).

After eliminating incorrect responses (2.5% of trials) and latencies <200 ms and >2000 ms (<1% of trials), we subjected the remaining latencies to an inverse ($1/x$) transformation (see Ratcliff, 1993). For interpretive ease, the means have been converted back to milliseconds. To control for the nested structure of the data, we employed multilevel analyses that simultaneously modeled variance at the lowest level (among different word types presented in different orders) and at the higher levels (among sentences and among participants; see Raudenbush & Bryk, 2002). Slopes were always treated as fixed effects; intercepts were allowed to vary randomly. Preliminary analyses revealed that trait, situation, and control words did not differ in the frequency of their distribution across the 8 positions in the lexical decision trials, $\chi^2(14, N=1008)=11.91$, $p=.61$, or the relative order in which they appeared during these trials, $\chi^2(1, N=510)=.20$, $p=.65$.

A 3 (response type) × 8 (lexical decision trial number) multilevel repeated-measures ANOVA on participants' transformed lexical decision latencies revealed a main effect of response type, $F(2, 726)=7.20$, $p<.001$. As illustrated in Fig. 1, relative to control words ($M=646$, $SD=84$), latencies were shorter for both trait words ($M=619$, $SD=75$), $F(1, 726)=6.33$, $p=.01$, and situation words ($M=613$, $SD=71$), $F(1, 726)=12.39$, $p<.001$. Latencies for trait and situation words did not differ, $F(1,726)=.75$, $p=.39$. This analysis also revealed a main effect of trial number, $F(7, 726)=4.24$, $p<.001$, indicating that latencies were generally longer on the first lexical decision trial immediately following the behavior description than on the remaining trials, which did not differ from one another. More importantly, the amount of time that had passed following the behavior description did not interact with response type, $F(14, 726)=1.35$, $p=.18$.

An additional analysis in which we submitted the transformed lexical decision latencies to a 2 (response type: trait vs. situation) × 2 (order: trait response first vs. situation response first) multilevel repeated-measures ANOVA revealed no main effects or interactions, $Fs(1, 248)<1.16$, $ps>.28$. Together, these analyses demonstrate that (a) on average, trait and situational inferences were activated to an equivalent degree, (b) this co-activation of trait and situational inferences remained constant with the passage of time (at least up to approximately 9 s later), and (c) the relative activation of trait versus situational inferences was not affected by which inference was assessed first.

Fig. 1. Mean lexical decision latencies for trait-relevant, situation-relevant, and matched-control words; error bars depict standard errors (Experiment 1).

Although these results are consistent with the co-occurring activation of both trait and situational inferences, it is possible that some participants primarily activated trait interpretations (and discounted the opposing situational interpretations), whereas others primarily activated situational interpretations (and discounted the opposing trait interpretations). Such a pattern of responding, when averaged across participants, could produce results resembling co-occurring activation. If this were the case, then participants' overall response times to trait words and situation words should be negatively correlated (Heider, 1958; Kelley, 1973). We found no support for this alternative explanation. Controlling for responses to control words to factor out overall differences in general reaction times, we observed a marginal *positive* correlation between responses to trait and situation words ($pr=.30$, $p=.10$).

Discussion

When incidentally exposed to behavior descriptions as secondary "distracter" stimuli within another task, people activated both trait and situational inferences, despite reporting no intention or awareness of having done so. These findings suggest that such inferences fulfill several criteria for automaticity (Bargh, 1994; Moors & De Houwer, 2006). Furthermore, participants' separate trait and situational inferences (a) did not differ in the time-course of their activation; (b) were not constrained by the previous assessment of other inferences; and (c) were, if anything, positively correlated with each other. Together, these latter findings suggest that the automatic inferences observed in this experiment were indeed co-occurring.

Experiment 2: Spontaneous, unconscious, and efficient co-occurring inferences

Experiment 2 was designed to extend Experiment 1 in several ways. First, we used a different experimental paradigm to examine automatic inferences. As part of a memory task, participants read a series of one-sentence behavior descriptions, after each of which they indicated whether a series of probe words had or had not appeared in the description. Second, we utilized a different control condition in which a separate group of participants responded to the same trait and situation probes following behavior descriptions that were semantically similar to those used in the experimental condition but did not afford the same inferences. The primary index of inference activation was the extent to which participants were slower to indicate that the trait and situation probe words had not actually appeared in the description following the inference-relevant versus the inference-irrelevant sentences (McKoon & Ratcliff, 1986). Furthermore, to examine whether the co-activation of multiple social inferences can occur effortlessly (as well as spontaneously and unconsciously), we had some participants complete a secondary task as they read the behavior descriptions.

Method

Participants

Participants were 91 native English-speaking students (55 women, 36 men) at a private university in the United States.

Procedure

As part of a study on "memory and linguistic processing," participants read a series of sentences, each one followed by a different set of words. Their objective was to indicate as quickly and accurately as possible whether each word had or had not appeared in the preceding sentence (i.e., probe-recognition task). Each block of probe-recognition trials began with a sentence presentation (3000 ms), followed by a fixation cross (500 ms), and then the first probe word. The probe remained on the screen until participants indicated whether or not the word had appeared in the preceding sentence by pressing one of two keys marked "Yes" and "No," respectively. Each trial was followed by a blank screen (250 ms), after which the next block of trials began. Lastly, participants were questioned about any explicit attempts (a) to figure out the causes of the behaviors, (b) to form an impression of the actor, or (c) to form an impression of the situation facing the actor.

Stimuli

In the *relevant probe* condition, the target sentences and the critical trait and situation words were taken from the lexical decision task in Experiment 1. To create an additional *irrelevant probe* condition, we altered these target sentences to produce 8 new sentences that included as many of the same words as possible, but did not afford the same trait and situational inferences (McKoon & Ratcliff, 1986). For example, "Nick avoids the big dog at the house on the corner" became "Nick and his dog walk by the big house on the corner." We then paired the critical trait ("afraid") and situation ("ferocious") probe words from the original sentences with these new semantically similar (but inferentially distinct) sentences. To further disguise the true purpose of the experiment, for each semantically similar pair of sentences, we included a *filler* probe word that also did not appear in either sentence. To equalize the number of "Yes" and "No" responses required, we also included 3 probe words that actually appeared in the sentence.

Cognitive load manipulation

Half the participants in each probe relevance condition were further assigned to a *cognitive load* condition. These participants completed a secondary task (see Todorov & Uleman, 2003) that involved rehearsing an 8-digit number, which appeared for 5 s before each block of probe-recognition trials (i.e., before the initial presentation of the behavior description). At the end of each block, participants completed a number-recognition trial, after which a new 8-digit number appeared and a new block of trials began. Changing the stimuli for the secondary digit-memorization task within each block ensured that cognitive load remained constant across the earlier and later trial blocks. The remaining participants in the *non-load* condition received neither the digit-presentations nor the number-recognition trials. Following the final block of trials, all participants rated the difficulty of the task (1 = *not at all difficult*, 7 = *very difficult*).

Results

Data from one participant were eliminated due to excessive errors (>50%) on the probe-recognition trials. No participant expressed any intention or awareness of having made inferences while reading the behavior descriptions; thus, any inferences observed can be surmised to be both spontaneous and unconscious.

Cognitive load manipulation check

An initial 2 (probe relevance) × 2 (cognitive load) ANOVA on the task difficulty ratings revealed that participants in the cognitive load condition ($M = 3.98$, $SD = 1.08$) perceived the task to be more difficult than did participants in the non-load condition ($M = 1.96$, $SD = 1.07$), $F(1, 87) = 78.77$, $p < .001$. No other effects approached significance ($Fs < 1$). Further analyses of the number-recognition trials within the cognitive load condition showed that, on average, participants responded correctly on 80% of the number-recognition trials, which is significantly higher than the chance level of 50%, $t(44) = 14.73$, $p < .001$, and this did not differ as a function of probe relevance ($F < 1$). Taken together, these results indicate that participants attended to the secondary digit-memorization task and that this increased attention resulted in a perceived depletion of the resources available for the primary task.

Because poor performance on the number-recognition task could indicate inattention to the secondary task, and thus an absence of cognitive load, we repeated all analyses reported as follows after removing participants who responded correctly on fewer than 75% of the number-recognition trials. These procedures did not change the significance of any of the results.

Analyses of inference activation

Response latencies. Participants' probe-recognition latencies served as our primary index of automatic inferences. Because inferences about traits and situations should interfere with probe-recognition judgments (i.e., make it more difficult to correctly respond "No" when presented with the relevant trait and situation words), longer latencies for trait and situation probes directly relevant to inferences afforded by the target sentences, relative to irrelevant trait and situation probes, would indicate the presence of such inferences. Furthermore, to the extent that the trait and situational inferences occur efficiently, the effect of probe relevance should not differ between the cognitive load and non-load conditions. Finally, if the activation of these inferences is truly co-occurring, then the effect of probe relevance should remain equivalent across both the time that passes between exposure to the descriptions and inference assessment, and the relative order in which trait and situation inferences are assessed.

After eliminating incorrect responses (1.3% of trials) and latencies <200 ms and >2000 ms (<1% of trials), we subjected the remaining latencies to an inverse ($1/x$) transformation (Ratcliff, 1993). For interpretive ease, the means have been converted back to milliseconds. We employed multilevel analyses in the same manner as Experiment 1. Preliminary analyses indicated that the pairings of the different types of response probes with the different cognitive load and probe-relevance conditions did not differ in their frequency of distribution across the 6 trials in the probe-recognition task, $\chi^2(39, N = 1441) = 46.17$, $p = .20$, or relative order, $\chi^2(11, N = 1441) = 12.09$, $p = .36$.

Submitting participants' transformed probe-recognition latencies to a 2 (probe relevance) × 2 (cognitive load) × 2 (probe type) × 6 (probe-recognition trial number) multilevel ANOVA, with repeated measures on the last two factors, first revealed that participants in the cognitive load condition ($M = 663$, $SD = 107$) responded more slowly overall than did participants in the non-load condition ($M = 623$, $SD = 111$), $F(1,87) = 5.03$, $p = .03$. This is further evidence for the effectiveness of the cognitive load manipulation. Critically, this analysis also revealed that participants who read behavior descriptions that afforded inferences relevant to the trait and situation probes ($M = 670$, $SD = 113$) responded to these probes more slowly than did participants who read behavior descriptions that were irrelevant to the trait and situation probes ($M = 617$, $SD = 102$), $F(1,87) = 7.39$, $p = .01$. Results further revealed that, in general, participants responded more slowly to situation probes than to trait probes, $F(1, 670) = 28.1$, $p < .001$, and that latencies were generally longer in the first probe-recognition trial following the behavior description than in any of the other positions, which did not differ from one another, $F(1, 670) = 13.92$, $p < .001$. Importantly, none of the higher-order interactions involving probe relevance was significant, $Fs(5, 670) < 1.46$, $ps > .20$.

An additional analysis in which we submitted the transformed probe-recognition latencies to a 2 (probe relevance) × 2 (cognitive load) × 2 (probe type) × 2 (probe order: first vs. second) multilevel ANOVA, with repeated measures on the last two factors, yielded the same main effects of probe relevance, cognitive load, and probe type as reported above, as well as a main effect of order, which indicated that responses were generally slower for whichever response came first, F(1,

$702) = 4.35$, $p = .04$. Again, none of the higher-order interactions involving probe relevance was significant, $Fs(1, 702) < 1.55$, $ps > .21$.

To more directly examine participants' separate trait and situational inferences, we conducted a 2 (probe relevance) × 2 (cognitive load) × 6 (trial number) multilevel ANOVA, with repeated measures on the last factor, for each type of probe. As displayed in Fig. 2, latencies for trait probes were longer following behavior descriptions for which the probes were relevant ($M = 656$, $SD = 108$) than following irrelevant descriptions ($M = 598$, $SD = 92$), $F(1, 87) = 8.27$, $p = .01$, an effect that was not moderated by cognitive load or trial number ($Fs < 1.12$). As also displayed in Fig. 2, latencies for situation probes were also longer following behavior descriptions for which the probes were relevant ($M = 683$, $SD = 118$) than following irrelevant descriptions ($M = 634$, $SD = 112$), $F(1, 87) = 4.76$, $p = .03$, an effect that again was not moderated by cognitive load or trial number ($Fs < 1.44$).

Lastly, to rule out the possibility that participants in the relevant probe condition simply responded more slowly in general than did participants in the irrelevant probe condition, we submitted latencies for the filler probes to a similar 2 × 2 × 6 multilevel mixed ANOVA. This analysis revealed the same cognitive load and trial number main effects described above. No simple or higher-order effects involving probe relevance emerged ($Fs < .40$, $ps > .61$).

Collectively, these results demonstrate that (a) on average, both trait and situational inferences were activated efficiently (i.e., even while under cognitive load) and to an equivalent degree, (b) their co-occurring activation did not differ depending upon when they were assessed (at least up to approximately 9 s later), and (c) their relative activation was not affected by which inference was assessed first. Because we assessed automatic inferences using a between-participants comparison of inference-relevant and inference-irrelevant target sentences, additional correlational analyses (as reported in Experiment 1) were not possible.

Error rates. In probe-recognition tasks, the interference caused by inference activation could result not only in slower correct identifications but also in a greater number of incorrect identifications (McKoon & Ratcliff, 1986). It is difficult to anticipate, however, whether participants will adopt a strategy that prioritizes speed concerns or accuracy concerns. As Uleman and colleagues note, "[i]t is not clear how to predict whether trait [or situation] inferences will affect reaction times or errors or both," and, to demonstrate the presence of automatic inferences, "[i]t is sufficient…to demonstrate either effect" (Uleman et al., 1996, p. 384). In the current experiment, parallel analyses of error rates showed no significant effects ($Fs < 1$). This is unsurprising given the very low error rates; nevertheless, it suggests that any observed differences in response times cannot be explained as an artifact of participants' sacrificing speed for accuracy in certain conditions.

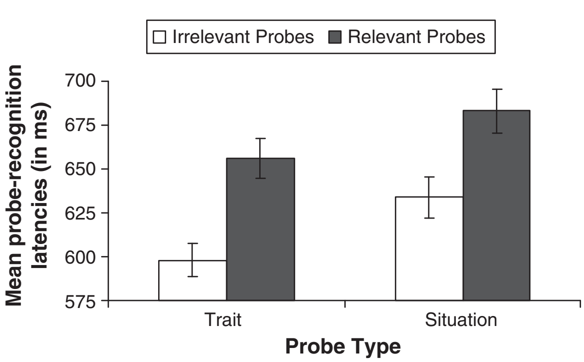

Fig. 2. Mean probe-recognition latencies for trait and situation probes as a function of probe relevance; error bars depict standard errors (Experiment 2).

Discussion

When processing behavior descriptions as part of a memory exercise, people automatically activated both traits and situational inferences, despite again reporting no awareness of having done so. Moreover, inference activation was not impaired by cognitive load, which fulfills yet another criterion of automaticity (i.e., efficiency; Bargh, 1994). Finally, participants' separate trait and situational inferences did not differ in the time-course of their activation, nor did their activation differ based on the relative order in which they were assessed, which is again consistent with the co-occurrence of these inferences. As discussed earlier, findings supporting the automatic and co-occurring activation of different types of social inferences appear to challenge some of the assumptions made by several accounts of the social inference process (Gilbert & Malone, 1995; Krull & Erickson, 1995a, but see Read & Miller, 1993, 1998; Trope, 1986).

Experiment 3: Spontaneous, efficient, and goal-independent co-occurring inferences

In Experiments 1 and 2, we examined incidental inferences (i.e., those formed when people encounter behaviors without any specific intention to make attributional judgments), which could be qualitatively different from inferences that are intentionally formed when deliberately thinking about others' behaviors (see Uleman, 1999). Perhaps when perceivers have specific goals to form impressions about actor or situations, these goals constrain what inferences will be automatically activated and what inferences require additional resources (Krull & Erickson, 1995a; Lieberman et al., 2002). Therefore, the primary objective of Experiment 3 was to investigate the co-occurring, automatic activation of multiple social inferences when people have a focal inference goal and only a limited amount of cognitive resources with which to pursue this goal. If inference goals and cognitive resources indeed constrain which social inferences people activate, then these two factors should moderate which types of inferences are activated and when this activation occurs. However, if the activation of multiple, co-occurring inferences is independent of perceivers' goals and cognitive resources, then these factors should have no additional influence on such activation.

The design of Experiment 3 was similar to Experiment 2, with several new variations. Impaired probe-recognition for trait and situation words directly relevant to the inferences implied by a target sentence again served as the primary index of inference activation. However, whereas Experiment 2 employed a between-participants manipulation of probe relevance, Experiment 3 manipulated this factor within participants. In addition, whereas participants in Experiments 1 and 2 read behavior descriptions as distracters or as part of a memory task, participants in Experiment 3 were explicitly instructed to form a trait impression or an impression of the situation.

Method

Participants

Participants were 93 native Dutch-speaking students (64 women, 29 men) at a private university in the Netherlands.

Procedure

As part of a "social judgment" task, participants read a series of different behavior descriptions. After reading each behavior description, they completed a word identification (i.e., probe-recognition) task. Each trial began with a row of Xs presented in the middle of the computer screen to signal the onset of the trial (1000 ms). Next, specific impression-formation instructions were displayed (3000 ms), followed by a one-sentence behavior description (3000 ms). A blank screen then appeared (500 ms), followed by another row of Xs (500 ms), and, finally, a probe word. The probe word remained on the screen until participants indicated whether or not the word had appeared in the preceding sentence by pressing one of two keys labeled "Yes" and "No," respectively. Once participants had responded, the next trial began. Thus, in this experiment, each behavior description was followed by a single probe word (vs. 8 in Experiment 1 and 6 in Experiment 2), thereby allowing us to include filler sentences in addition to the critical sentences.

Stimuli

The behavior descriptions and probe words for this experiment were, for the most part, Dutch versions of the descriptions and probe words from Experiments 1 and 2; we selected six sets of descriptions and probes from Ham and Vonk (2003). For the relevant probe trials, each behavior description was followed by a probe word that was implied (as determined by pre-testing, see Ham & Vonk, 2003), but never explicitly mentioned, in the behavior description (e.g., "Don can't get the computer started" was followed by "inept"). The irrelevant probe trials consisted of the same behavior descriptions and probe words as the relevant probe trials, but the probes followed a behavior description for which it was irrelevant (e.g., "Will talks during the lecture" was followed by "inept"). In addition to the 12 relevant probe (6 traits, 6 situations) and 12 irrelevant probe (6 traits, 6 situations) trials paired with the critical behavior descriptions, we also included 36 filler trials consisting of different behavior descriptions. These filler trials helped to disguise which trials were of critical interest. Of these filler trials, 6 featured a probe word that did not appear in the sentence (as was the case with the 24 critical behavior descriptions) and 30 featured a probe word that actually appeared in the sentence, thereby ensuring that "Yes" and "No" responses across the entire set of 60 trials were required with equal frequency.

Inference goal manipulation

We manipulated inference goals by providing participants with both general instructions in the introduction to the experiment, and short, specific instructions before each of the 60 behavior descriptions. In the *trait goal* condition, participants were initially informed about the influence of personality traits on behavior. Then, before reading each behavior description, they were asked to "form an impression of [the actor in the description]." In the *situation goal* condition, participants were initially informed about the influence of situational factors on behavior. Then, before each behavior description, they were asked to "form an impression of [the specific situation outlined by the description]." For example, for the behavior, "Don can't get the computer started," participants in the trait goal condition were asked to form an impression of *Don*, whereas participants in the situation goal condition were asked to form an impression of *the computer*. Although inferences that are congruent with the assigned inference goal (e.g., trait inferences following a trait inference goal) are clearly not automatic, any additional inferences that are also elicited can be considered automatic.

Cognitive load manipulation

Participants in the *cognitive load* condition completed a secondary memorization task that differed from the one used in Experiment 2. Gilbert, Krull, and colleagues have shown that when perceivers have a specific goal of inferring either traits or situations, cognitive load prevents the integration of alternative inferences, even when the cognitive load is itself created by rehearsing information directly relevant to such alternatives (Gilbert et al., 1988; Krull & Erickson, 1995b). In the present experiment, one word in each behavior description was underlined. Participants in the *cognitive load* condition were instructed to memorize the underlined words. For each description, either the actor's name or the situation the actor was experiencing (e.g., "Don can't get the computer started" or "Don can't get the computer started") was underlined. Whether the trait-underlined or situation-underlined version of each sentence appeared was determined on a random basis for each participant. After every 6 trials, we asked participants to list all six underlined words. Participants in the *non-load* condition were asked to ignore this underlining.[2]

As in previous studies by Gilbert, Krull, and colleagues (Gilbert et al., 1988; Krull & Erickson, 1995b), the cognitive load manipulation itself draws attention to information that implies the possibility of alternative interpretations; nevertheless, participants must still generate these interpretations (i.e., that Don is inept or that the computer is broken) themselves. Thus, despite the salience of trait and situational information created by the cognitive load manipulation, if participants exhibit activation of both trait and situational inferences in this condition, this is still evidence for automatic inferences that are not constrained by the presence of specific inference goals.

Results

Cognitive load manipulation check

Analyses of the correctly recalled underlined words in the cognitive load condition revealed adequate performance on this task. On average, participants correctly recalled 78% of the underlined words, and this did not differ as a function of inference goal ($F<1$). This indicates that participants were indeed attending to the secondary word-memorization task. As in Experiment 2, repeating the analyses below after eliminating participants who correctly recalled fewer than 75% of the underlined words did not change the significance of any of the results.

Analyses of inference activation

Response latencies. As in Experiment 2, participants' probe-recognition latencies served as our primary index of automatic inferences. To the extent that participants activate both trait and situational inferences, even in the presence of specific impression-formation goals, the effect of probe relevance should not differ by probe type, inference goal, or any interaction involving these two factors. Moreover, to the extent that trait and situational inferences are also efficient, the effect of probe relevance should not differ between the cognitive load and non-load conditions. Unlike in Experiments 1 and 2, participants responded to a single probe after each behavior description; consequently, it was not possible to examine the specific time-course of inference activation. It was possible, however, to examine whether the presentation order of the relevant and irrelevant trait and situation probes affected inference activation. As in previous experiments, if inference activation is truly co-occurring, then the effect of probe relevance should be unaffected by the relative order in which trait and situation responses were assessed.

After eliminating incorrect responses (2.1% of trials) and responses <200 ms and >2000 ms (<2% of trials), we subjected the remaining responses to an inverse ($1/x$) transformation (Ratcliff, 1993). For interpretive ease, the means have again been converted back to milliseconds. We employed multilevel analyses as in Experiments 1 and 2. Preliminary analyses indicated that the pairings of the different types of response probes with the different inference goal and cognitive load conditions did not differ in their frequency of distribution across the 4 possible orders in the probe-recognition task, $\chi^2(56, N=2143)=69.29$, $p=.11$.

[2] To determine any additional effects that underlining the name of the actor versus the situation might have had (beyond being part of the cognitive load manipulation), we repeated all analyses with this additional factor included. The particular portion of the sentence that was underlined had no significant simple or higher-order effects either in this experiment or in Experiment 4 (all Fs<1).

Submitting participants' transformed probe-recognition latencies to a 2 (inference goal) × 2 (cognitive load) × 2 (probe relevance) × 2 (probe type) × 4 (order of probe presentation) multilevel ANOVA, with repeated measures on the last three factors, revealed that participants in the cognitive load condition (M = 814, SD = 157) responded more slowly overall than did participants in the non-load condition (M = 702, SD = 166), F(1, 89) = 12.35, p < .001. This suggests that the alternative secondary memory task used in this experiment was effective. Critically, latencies for trait and situation probe words that were directly relevant to the inferences afforded by the behavior descriptions (M = 810, SD = 179) were longer than latencies for trait and situation probe words that were irrelevant to such inferences (M = 768, SD = 162), F(1, 1526) = 33.64, p < .001, an effect that was not further moderated by any higher-order interactions involving probe type, inference goal, cognitive load, or order (Fs < 2.66, ps > .10). There was also a main effect of order, F(3, 1526) = 12.13, p < .001, indicating that latencies were generally longer for whichever probe was presented first or second as compared to the probe presented third, which had longer latencies than whichever probe was presented fourth.

To more directly examine participants' separate trait and situational inferences, we conducted a 2 (inference goal) × 2 (cognitive load) × 2 (probe relevance) × 4 (presentation order) multilevel ANOVA, with repeated measures on the last two factors, for each type of probe. As displayed in Fig. 3, latencies for relevant trait probes (M = 807, SD = 179) were longer than latencies for irrelevant trait probes (M = 768, SD = 164), F(1, 493) = 14.01, p < .001. Similar to the overall analysis presented above, higher-order interactions involving probe relevance, cognitive load, and inference goals were not significant (Fs < 2.5, ps > .12). As also displayed in Fig. 3, latencies for relevant situation probes (M = 813, SD = 195) were longer than latencies for irrelevant situation probes (M = 768, SD = 172), F(1, 483) = 19.07, p < .001. Once again, higher-order interactions involving probe relevance, cognitive load, and inference goals did not approach significance (Fs < 1.5, ps > .23). Together, these analyses again demonstrate that, on average, both trait and situational inferences were activated efficiently (i.e., even while under cognitive load) and that this activation was unconstrained by a conscious focus on trait versus situational interpretations of behavior. In addition, the relative activation of trait versus situational inferences was unaffected by the order in which these inferences were assessed.

As in Experiment 1, all participants responded to both inference-relevant and inference-irrelevant probes. It was therefore again possible to test whether separate subsets of participants showed a relative preference for trait inferences and discounted situational inferences, whereas others showed a relative preference for situational inferences and discounted trait inferences, a pattern of responding that would resemble co-occurring activation when averaged across participants. To examine this alternative explanation for the current results, we created a trait inference index by subtracting latencies for irrelevant trait probes from latencies for relevant trait probes and a situational inference index by subtracting latencies for irrelevant situation probes from latencies for relevant situation probes. If some participants are primarily activating trait inferences and discounting the opposing situational inferences or vice-versa, these indices should be negatively correlated (Heider, 1958; Kelley, 1973). Controlling for filler probe latencies to factor out individual differences in general response speed, results revealed that responses to trait and situational probes were uncorrelated (pr = .02, p = .85), suggesting once again that participants are not primarily activating one inference and discounting the other.

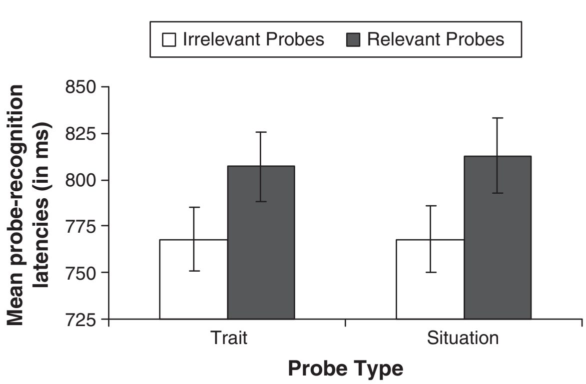

Fig. 3. Mean probe-recognition latencies for trait and situation probes as a function of probe relevance; error bars depict standard errors (Experiment 3).

Error rates. Although error rates were again extremely low, parallel analyses conceptually replicated the results for response latencies. Following a square root transformation (Cohen & Cohen, 1975), we submitted participants' errors to a similar 2 × 2 × 2 × 2 × 4 ANOVA. For interpretive ease, the means have been converted back to raw errors. This analysis revealed only a main effect of probe relevance, F(31, 1614) = 24.16, p < .001. False recognition was higher for trait and situation probes that were directly relevant to the inferences afforded by the behavior descriptions (M = .54, SD = 1.40) than for irrelevant trait and situation probes (M = .14, SD = 0.38). There were no higher-order interactions involving probe type, inference goal, cognitive load, or order (Fs < 1.22, p > .27), which further suggests that participants efficiently formed multiple, co-occurring social inferences that were unaffected by their explicit inference goals.

Discussion

Even when participants deliberately pursued either a trait or a situational inference goal, they still activated both types of inferences, regardless of cognitive load (i.e., there was no moderation of this activation by inference goal or by cognitive load). In addition, (a) these separate trait and situational inferences were not constrained by the order in which they were assessed, and (b) the tendency to form one type of inference was independent of the tendency to form the other type of inference. These results, which replicate and extend the findings of Experiments 1 and 2, more clearly indicate that, contrary to the assumptions of several current models of social inference (Gilbert & Malone, 1995; Krull & Erickson, 1995a; Lieberman et al., 2002), perceivers spontaneously and efficiently activate co-occurring trait and situational inferences, even while pursuing an impression-formation goal that specifically targets only one of these inferences.

Taken together, the results of Experiments 1–3 provide strong evidence that when encountering (written descriptions of) others' behaviors, perceivers interpret them in several different ways and with a high degree of automaticity that is not constrained by specific impression-formation goals. As noted earlier, given these findings, it becomes difficult to explain the robust biases that have long been observed in attributional judgments (Gilbert & Malone, 1995; Molden & Dweck, 2006; Ross & Nisbett, 1991) in terms of the unequal activation of different interpretations (see also Gawronski, 2004).

How then do such biases arise? As noted earlier, Ham and Vonk (2003) have speculated that the initial, efficient interpretations described in prior research (Gilbert et al., 1988; Krull & Erickson, 1995b; Lee & Hallahan, 2001) may actually reflect a secondary process of *selection* among multiple, automatically activated inferences. That is, although multiple inferences may initially be automatically activated, one of these inferences may subsequently be assigned greater weight or deemed more relevant or valid (Read & Miller, 1993, 1998; cf. Gawronski & Bodenhausen, 2006). According to this view, instead of preventing the activation of alternative interpretations, the greater contextual salience (i.e., diagnosticity or perceptual prominence) of particular interpretations or the presence of specific inference goals could later increase the likelihood that one active interpretation achieves greater prominence in (i.e., is "selected for") perceivers' final impressions through one or more of these mechanisms.

One basic implication of this type of selection mechanism is that directly assessing people's deliberate attributional judgments (e.g., via self-report) of the same behaviors that elicited co-occurring activation of multiple inferences in Experiments 1–3 should reveal the typical effects described by traditional models of social inference (Gilbert & Malone, 1995; Krull & Erickson, 1995a; Lieberman et al., 2002). That is, if some secondary selection process occurs, then people's explicit explanations should be biased in line with the inference goal they are pursuing, particularly when they do not have the cognitive capacity to overcome these biases by integrating alternative attributions. This pattern would produce an interaction between participants' inference goals and cognitive load. Such findings would provide preliminary evidence for a distinction between the automatic activation of multiple inferences when initially exposed to a particular behavior (as assessed in Experiments 1–3) and the more deliberate selection of one of these inferences when explicitly forming impressions of this same behavior.

Experiment 4: Deliberate attributional judgments

The materials and manipulations used in Experiment 4 were nearly identical to those from Experiment 3. The one critical difference was that instead of assessing automatic inference activation, we directly assessed trait and situational attributions using explicit reports. As noted above, to the extent that such attributional judgments reflect a selection among a set of activated inferences rather than a simple report of those inferences that are active, then (a) these judgments should be biased by the inference goals that participants are pursuing, and (b) these biases should be most (and perhaps only) apparent when experiencing cognitive load (Gilbert & Malone, 1995; Krull & Erickson, 1995a; Lieberman et al., 2002). That is, as demonstrated by Gilbert, Krull, and colleagues (Gilbert et al., 1988; Krull, 1993; Krull & Erickson, 1995b; see also Lee & Hallahan, 2001; Molden et al., 2006), participants' attributions should be skewed in the direction of their inference goals when they are under cognitive load and thus unable to effortfully consider additional interpretations that do not match their focal goals. However, as also demonstrated by previous research, when participants are not under cognitive load, their attributions should not be skewed by their inference goals because they will have the cognitive capacity to consider multiple interpretations for the behaviors.

Method

Participants

Participants were 97 native Dutch-speaking students (66 women, 31 men) at a private university in the Netherlands.

Procedure

The procedure and materials were very similar to those used in Experiment 3. Participants received specific instructions to form either trait or situational inferences, after which a one-sentence behavior description appeared. Participants in the cognitive load condition were further asked to memorize underlined words relating either to the actor's name or the situation he was experiencing; participants in the non-load condition were instructed to ignore the underlining. The behavior descriptions were identical to those used in Experiment 3. Whereas in Experiment 3 it was necessary to present the 6 critical descriptions multiple times during the probe-recognition task, it was only necessary to present each behavior once in this experiment.

To capture deliberate trait and situational attributions, after each behavior description, we asked participants to rate the extent to which the actor possessed the trait implied by the behavior and the extent to which the actor's situation might have influenced his behavior. For example, for the behavior, "Will talks during the lecture," participants answered the question, "How impolite is Will?" (1 = *not at all impolite*, 7 = *very impolite*), and the question, "How boring is the lecture?" (1 = *not at all boring*, 7 = *very boring*). For each description, trait and situation questions were presented in random order.

Results

Manipulation checks

Analyses of the correctly recalled underlined words showed that performance was again adequate. On average, participants correctly recalled 83% of the underlined words, and this did not differ as a function of inference goal ($F<1$), suggesting that participants in the cognitive load condition attended to the secondary task. As in previous experiments, repeating the analyses reported below after removing participants who correctly recalled fewer than 75% of the underlined words did not change the significance of any of the results.

To provide an additional check on the cognitive load manipulation and to determine the effectiveness of the inference goal manipulation, we examined the time it took participants to report their attributional judgments. A 2 (inference goal) × 2 (cognitive load) × 2 (attribution type) mixed ANOVA revealed that participants took longer to answer the attributional questions in the cognitive load condition than in the non-load condition, $F(1, 93) = 15.11$, $p<.001$. This analysis also revealed an inference goal by attribution type interaction, $F(1, 93) = 43.85$, $p<.001$, indicating that participants answered attributional questions that were consistent with their focal inference goal (e.g., trait attributions following a trait goal) more quickly than they answered questions that were inconsistent with this goal (e.g., situational attributions following a trait goal).

That goal-consistent judgments were made more easily than goal-inconsistent judgments suggests that participants adopted the appropriate inference goal while processing the actors' behaviors (see Krull, 1993). No other effects approached significance ($F\text{s}<1.18$, $p\text{s}>.27$). Together, these data confirm the effectiveness of both the cognitive load and inference goal manipulations.

Deliberate attributional judgments

To the extent that attributions reflect a biased selection among multiple activated inferences, one would expect stronger trait attributions among participants pursuing trait inference goals and stronger situational attributions among participants pursuing situational inference goals. Furthermore, the influence of inference goals should be more pronounced when participants' cognitive resources are depleted (i.e., when they are unlikely to be able to consider alternative interpretations; see Gilbert et al., 1988; Krull & Erickson, 1995b). Overall, these effects should produce an inference goal × cognitive load × attribution type interaction.

Submitting participants' attributional judgments to a 2 (inference goal) × 2 (cognitive load) × 2 (attributions type) mixed ANOVA, with repeated measures on the last factor, revealed that, in general, participants formed stronger trait attributions ($M = 5.12$, $SD = .81$) than situational attributions ($M = 4.51$, $SD = .94$), $F(1, 93) = 7.26$, $p<.01$. This analysis also revealed both the predicted inference goal x attribution type interaction, $F(1, 93) = 27.60$, $p<.001$, as well as the predicted cognitive load × inference goal × attribution type interaction, $F(1, 93) = 33.60$, $p<.01$.

To better understand this three-way interaction, we separately examined attributional judgments in the cognitive load and non-load conditions. As displayed in Fig. 4, participants in the non-load condition did not show any differences in their attributional judgments, regardless of their inference goal (all $Fs < 1$). This suggests that, consistent with previous research (Gilbert et al., 1988; Krull, 1993; Krull & Erickson, 1995b; Lee & Hallahan, 2001; Molden et al., 2006), when participants had ample cognitive resources, they integrated both trait and situation interpretations into their attributional judgments. In the cognitive load condition, however, a significant inference goal × attribution type interaction emerged, $F(1, 94) = 91.21$, $p < .001$. As also displayed in Fig. 4, participants reported stronger trait attributions when pursuing trait goals ($M = 5.49$, $SD = .67$) versus situational goals ($M = 4.45$, $SD = .83$), $F(1, 94) = 28.83$, $p < .001$, and stronger situational attributions when pursuing situational goals ($M = 5.39$, $SD = .76$) versus trait goals ($M = 3.92$, $SD = .65$), $F(1, 94) = 58.65$, $p < .001$.

Replicating previous work on trait and situational attributions, inference goals biased attributional judgments, but only when participants' attentional resources were depleted. Moreover, these biases emerged for a set of behaviors that were known (based on the results of Experiment 3) to elicit the automatic and co-occurring activation of both trait and situational inferences, a finding that is consistent with our proposition that such attributions reflect a selection among activated inferences. Correlational data lend further support to this selection mechanism. Whereas the correlation between trait and situational attributions (averaged across the 6 target behaviors) was negative and significant in the cognitive load condition ($r = -.37$, $p = .003$), this correlation was positive and non-significant in the non-load condition ($r = .09$, $p = .61$); these correlations were significantly different from each other ($Z = 2.15$, $p = .031$). Thus, unlike in Experiment 3, participants were primarily selecting either trait or situational interpretations and discounting the opposing alternative (Heider, 1958; Kelley, 1973), but only when their cognitive resources were in short supply, leaving them unable to integrate the alternative interpretation.

Discussion

When asked to provide deliberate attributional judgments of the exact same behaviors that had been shown in Experiment 3 to elicit the automatic and co-occurring activation of trait and situational inferences, participants' responses were indeed affected by both their specific inference goals and whether or not they were experiencing cognitive load. That is, although inference goals and cognitive load did not affect inference activation in Experiment 3, these factors exerted a pronounced effect on attributional judgments. Thus, such judgments do not appear to reflect a simple reporting of whatever inferences have been activated; rather, they may reflect the outcome of a selection process, wherein one of these activated inferences is chosen and the other is discounted.

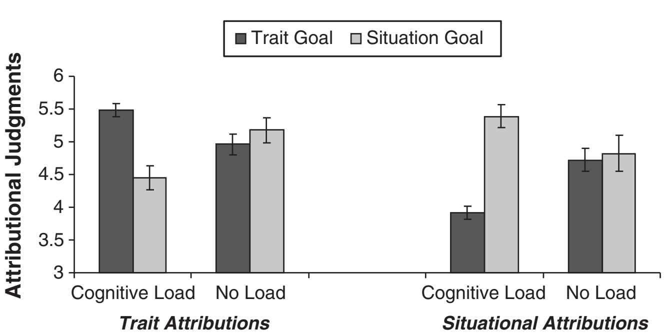

Fig. 4. Mean trait and situation attributions as a function of inference goal and cognitive load; error bars depict standard errors (Experiment 4).

General discussion

Although our social perceptions feel as immediate and reliable as our perceptions of the physical world, the current experiments add to a long history of research demonstrating that social perception is a complex, multifaceted process that can be biased by a variety of factors. Where these experiments differ from previous research, however, is in the evidence they provide for the source of biases in social inference and how these biases arise.

The current research demonstrates that upon encountering others' behaviors, perceivers activate multiple, co-occurring inferences (a) in the absence of intention and awareness (Experiments 1 and 2), (b) in the presence of cognitive load (Experiments 2 and 3), and (c) independent of specific inference goals (Experiment 3). Moreover, because in several experiments (Experiments 2 and 3) making social inferences actually interfered with what participants believed was their primary task (i.e., rapid and accurate recognition of probe words), there is at least suggestive evidence that these inferences are not easily controlled (see also Ham & Vonk, 2003; Uleman et al., 1996). These results extend the findings of Ham and Vonk (2003) by showing not only that perceivers can spontaneously activate both trait and situational inferences, but also that this activation is unconscious, efficient and co-occurring. Furthermore, we replicated these findings using (a) two different measures of automatic inference (facilitation in a lexical decision task and interference in a probe-recognition task), (b) experimental procedures featuring both within-persons and between-persons designs, and (c) participants from two different countries (the United States and the Netherlands). On the whole, the automatic and co-occurring activation of multiple social inferences appears to be a robust phenomenon.

Despite the robustness of our results, some qualifications are warranted. First, one drawback of the probe-recognition task used in Experiments 2 and 3 is that it does not fully discriminate between on-line inferences made during encoding (i.e., while processing the behavior descriptions) versus those made retrospectively when testing recognition (Keenan, Potts, Golding, & Jennings, 1990). Although some influence of retrospective recognition processes cannot be completely ruled out, research has shown that these processes are limited when (a) test probes are presented directly after encoding, and (b) participants are experiencing cognitive load (McKoon & Ratcliff, 1986). Both conditions were met in Experiments 2 and 3, and the absence of any effects of cognitive load on inference activation further suggests that these effects were not entirely a product of recognition processes rather than online inference. Moreover, retroactive recognition processes are a highly unlikely alternative explanation for the results of Experiment 1, which measured inference activation in terms of facilitated responses during a lexical decision task (see Keenan et al., 1990).

Second, much of the evidence supporting the automaticity of trait and situational inferences was based on the absence of moderation by particular factors – inference goals and cognitive load – previously shown to affect deliberate attributional judgments (e.g., Gilbert et al., 1988; Krull, 1993; Krull & Erickson, 1995b; Lee & Hallahan, 2001; Molden et al., 2006). Although there is always some uncertainty in the conclusions that can be drawn from the specific absence of particular effects, it is important to note that our cognitive load manipulations significantly altered overall response times in Experiments 2–4 and that our inference goal manipulation altered deliberate attributional judgments in Experiment 4. Thus, we are not basing our arguments for automaticity on manipulations that produced null effects; instead, as is common in the automaticity literature (see Moors & De Houwer, 2006), we are inferring automaticity from the specific lack of effects of these manipulations on the activation of trait and situational inferences, despite the clear effects they had on other measures.

Third, the current research relied solely upon responses to written behavior descriptions. These descriptions might instigate processing that is qualitatively different from processing that occurs when observing behaviors *in vivo*. In other words, our results may be more relevant for the general comprehension of written behavioral descriptions than for the immediate perception of live social behaviors. However, Fiedler and Schenck (2001) have documented that automatic trait inferences are elicited by pictorially presented behaviors in a way that is analogous to their elicitation by verbal descriptions. Also, Malle (2008) has recently performed direct comparisons of the effects of verbal descriptions and visual presentations of the same behaviors and has found no differences as a function of presentation mode in the extent to which inferences about traits, goals, beliefs, and intentions are co-activated. Although further research is needed to determine whether the present results hold when perceivers actually observe others' actions, evidence exists suggesting that our findings would generalize to these circumstances. Even if the automatic and co-occurring activation of multiple inferences is limited to verbal behavior descriptions, this is still a primary medium through which people learn about their social world, and thus our findings still have important implications for understanding social perception processes.

Finally, although many types of inferences are elicited by others' behaviors (Malle, 1999; Read & Miller, 1993, 1998; Reeder, 2009), our experiments focused exclusively on the co-occurrence of trait and situational inferences. As noted earlier, this approach provided the strongest test of the possibility of automatic and co-occurring inferences, in that judgments about traits and situations have typically been thought to operate in a "hydraulic" or opposing manner (Gilbert & Malone, 1995; Heider, 1958; Kelley, 1973; Ross & Nisbett, 1991). Furthermore, because trait and situational inferences lie at the heart of prominent models of social inference (e.g., Gilbert et al., 1988; Krull & Erickson, 1995a; Trope & Alfieri, 1997; Trope & Gaunt, 2000), focusing on these inferences also ensured that the current results would be directly comparable to this previous research. We should reiterate, however, that although trait and situational inferences have important implications for social perception, we do not claim that these are the only important inferences that perceivers draw. Instead, we are arguing that the automatic co-occurrence of even somewhat incompatible inferences suggests that other more compatible inferences – for instance, those concerning intentions and desires – are also likely to co-occur.

Toward an expanded account of automatic processes in social inference

The mounting empirical evidence for the automatic activation of multiple social inferences reported here and elsewhere (Ham & Vonk, 2003; see also Reeder, Vonk, Ronk, Ham, & Lawrence, 2004) poses a challenge to several traditional accounts of social inference, which have typically characterized the consideration of multiple interpretations of behavior as an effortful, resource-dependent process and which have discussed a variety of constraints (e.g., perceivers' specific inference goals) on whether trait or situation interpretations are automatically activated (Gilbert, & Malone, 1995; Krull & Erickson, 1995a; Lieberman et al., 2002). Although biases are clearly evident in deliberate attributional judgments when people have specific inference goals and are cognitively busy (Gilbert et al., 1988; Krull & Erickson, 1995b; Lee & Hallahan, 2001; Molden et al., 2006), our findings indicate that such biases do not stem from a failure to activate alternative interpretations of behavior (see also Gawronski, 2004). Instead, they suggest that inferential biases and inference activation may occur independently of each other.

If biases in social inference do not depend upon the constrained activation of particular interpretations, then what process(es) might account for their emergence? We propose that many of the commonly observed social inference biases stem from some type of "selection" among multiple co-active interpretations (see also Ham & Vonk, 2003). That is, people's initial interpretations may be better understood as a rapid (and perhaps implicit) "choice" among several accessible interpretations that is guided by their dominant impression-formation goals (e.g., Gilbert et al., 1988; Krull & Erickson, 1995b), by their lay theories about the causes of behavior (e.g., Knowles, Morris, Chiu, & Hong, 2001; Molden et al., 2006), or by additional contextual information that boosts the salience of particular interpretations (e.g., Trope & Gaunt, 2000).

Although Experiment 4 provides preliminary support for selection processes, this evidence is indirect and stems from quasi-experimental procedures. That is, we assume, but did not directly demonstrate, that the behavior descriptions actually elicited the automatic and co-occurring activation of trait and situational inferences just as they had in Experiment 3. Additional research is clearly needed to provide more concrete support for selection processes (e.g., evidence for selection and evidence for co-occurring activation of multiple inferences within a single experiment) and to define the specific mechanisms through which these processes might operate (e.g., whether selection takes place in sequence or in parallel with the automatic activation of multiple inferences, or whether selection itself can be a spontaneous and efficient process). Nevertheless, given the evidence for the automatic and co-occurring activation of multiple inferences, it is useful to consider several potential mechanisms outlined by previous research that are broadly consistent with our selection proposal.

Selection as enhanced or sustained activation of particular inferences

Although multiple inferences may be activated when observing a particular behavior, the activation of certain inferences could be selectively enhanced or sustained for a longer period of time based upon perceivers' lay theories or other information about the context in which the behavior itself occurs. That is, although many applicable inferences may become active compared to some resting baseline level, a selection among these inferences may be further primed by certain characteristics of the environment in which a behavior is observed or of the specific person who has observed it. If this type of priming were to occur, then inferences with the highest current activation should win the "race" to the threshold of awareness and should "capture" (i.e., dominate or bias) the final impression that is formed about the behavior (cf. Higgins, 1996). In some circumstances, whichever inference is selected in this manner subsequently may even actively inhibit alternative inferences (Read & Miller, 1993, 1998), which could further explain why overcoming social inference biases can require careful attention or substantial motivation (Gilbert & Malone, 1995; Ross & Nisbett, 1991). The idea of selection as the enhanced or sustained activation of particular inferences bears a resemblance to the parallel constraint satisfaction processes that underlie Read and Miller's (1993, 1998) model of social perception, and is consistent with the results of Experiment 4. Future research could directly investigate this type of selection process by first assessing initial levels of automatic inference activation and then assessing this activation again later following an explicit impression-formation task.

Selection as the "attachment" of particular inferences to actors

Aside from a quantitative difference in the relative accessibility of different inferences, selection mechanisms could involve a combination of qualitatively different processes. For example, Todorov and Uleman (2002, 2003, 2004) have made a distinction between the *activation* of trait inferences and the *attachment* (or *binding*) of these inferences to actors. Their findings show that when initially presented with photographs of individuals and descriptions of behavior ostensibly performed by those individuals, people falsely remember having seen a trait word implied by the behavior as having been presented with that photograph. This effect persists even when participants view two photographs with each description but are told that the behavior only refers to one of the individuals depicted

(Todorov & Uleman, 2004). That is, trait concepts appear to be selectively attached only to the actors whom they are purported to describe, not those individuals who happen to be present when the trait inferences are activated (see also Brown & Bassili, 2002). Thus, as with the present experiments, the mere activation of trait concepts in the presence of certain individuals was not sufficient to bias participants' later responses; such biases required some form of additional selection process that further bound these trait concepts to a particular actor.

Although Todorov and Uleman discuss the process of trait inference attachment, this idea of an additional representational link between actors and possible behavioral interpretations can be extended to other inferences about the goals, intentions, or beliefs that a particular action implies. Moreover, the concept of attachment could perhaps even be extended to situational inferences. Consistent with this possibility, Kammrath, Mendoza-Denton, and Mischel (2005) have demonstrated that, when people explicitly attempt to form trait judgments, they can incorporate information about the context surrounding an actor's actions in terms of a conditional *if…then* impression (e.g., “if pressured, then she is obnoxious”). Automatically activated situational inferences could thus possibly be attached to actors through similar conditional representations (e.g., “when a ferocious dog was present, Nick avoided it”). However, given this less direct means of attachment, it may be that situational inferences are less likely to be or are more weakly linked to actors than are trait inferences. These circumstances might help to explain why, even though trait and situation interpretations are equally likely to be activated from observations of behavior (as the current work suggests), biases favoring trait judgments are, at times, more prevalent.

Selection as attribution versus mere association

Carlston and Skowronski (2005) have made a similar qualitative distinction between inference activation and the utilization of these inferences in attributional judgments. They have articulated two separate processes that can follow the activation of trait information: (a) the mere *association* of an activated trait inference with a bystander who is present when this activation occurs, and (b) the more elaborate *attribution* of this trait to a particular actor. Their studies have demonstrated that when one person is describing the behavior of another, the trait inferences activated by this description are typically associated with both the communicator and the perpetrator (see also Mae, Carlston, & Skowronski, 1999; Skowronski, Carlston, Mae, & Crawford, 1998). Furthermore, such associations, they argue, are relatively automatic and are not altered by cognitive load, different processing goals, or even explicit warnings that these associations are not accurate. In contrast, when a person is describing his or her own behavior, the trait inferences that are activated typically result in more elaborate attributions about this person. Such attributions are more susceptible to disruption by cognitive load, are more dependent on particular impression-formation goals, and are eliminated by warnings that these attributions are not true.

Thus, Carlston and Skowronski's (2005) account also draws a distinction between the immediate consequences of inference activation, in terms of the associations that form, and the consequences of a secondary selection process by which active inferences are definitively attributed as causes of the behavior. Furthermore, although from their perspective the biases that are most typically observed in social inference research could most readily be conceptualized as resulting from the selection of a particular attribution from among those that had been automatically activated, this perspective also illuminates the possibility of more subtle biases that could arise from the mere association of other automatically activated inferences with others involved with this behavior (see also Goren & Todorov, 2009).

In summary, there is evidence supporting a variety of mechanisms through which biases in social inference could arise following the co-occurring, automatic activation of multiple possible interpretations. Each of these mechanisms involves what we have broadly characterized as “selection” processes. Additional research on the distinction between activation and selection processes could provide new insights into how the impressions triggered by others' actions unfold and why they are often biased in systematic ways.

Summary and conclusions

The current research provided evidence for the automatic and co-occurring activation of multiple social inferences. Based upon this evidence, we have taken steps toward articulating an expanded account of how social inference biases arise, one that involves an additional process of selection among the multiple inferences that have been activated. On the surface, the difference between the present account and traditional perspectives on social inference may not seem great. However, this difference has several important implications. First, two prominent explanations for the robust tendency to draw correspondent trait inferences have been that social perceivers fail either (a) to notice or (b) to appreciate the relevance of other possible interpretations (particularly those involving the social situations in which the behaviors occur; see Gilbert & Malone, 1995; Ross & Nisbett, 1991). However, the current research complements previous research arguing that people do recognize and consider other alternatives beyond traits when observing behavior (Gawronski, 2004; Malle, 1999; Read & Miller, 1993, 1998; Reeder, 2009). Our expanded account therefore implies that other explanations for the prominence of trait-focused biases – ones that involve how and when alternate explanations are integrated (Krull & Erickson, 1995a; Trope & Gaunt, 2000) – may more frequently account for this phenomenon (Gawronski, 2004).

Second, if multiple interpretations are automatically activated, then, even if one interpretation is selected or preferred, the other active inferences may still exert some residual influence on attributional judgments. Just as Carlston and Skowronski (2005) have demonstrated that activated trait interpretations can have measurable associations with an individual even if these traits are not directly attributed to that person (see also Brown & Bassili, 2002; Goren & Todorov, 2009), so too may the multiple interpretations that are automatically activated when encountering behavior continue to have similar associative effects even if they are not completely integrated into the social impressions that people deliberately form. As noted above, our expanded account of social inference processes therefore also implies that, in addition to the prominent biases that are typically observed in social inference, there may be additional, more subtle biases that have not yet been fully identified but that might play an important role in people's social impressions.

To conclude, the current research suggests that upon encountering a hapless colleague and his recalcitrant cell phone, you would likely automatically activate a variety of inferences about his character, the situational pressures he is facing, as well as his goals and intentions. The challenge for future research is to further unravel the processes by which you might select among these different interpretations when attempting to integrate them into a coherent social impression and to more thoroughly understand the ways in which bias might emerge in the impression that you ultimately form.

References

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. WyerJr., & T. K. Srull (Eds.), *Handbook of social cognition (Vol. 1)* (pp. 1−40)., 2nd Ed. Mahwah, NJ: Lawrence Erlbaum Associates.

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.

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.

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.

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

Cohen, J., & Cohen, P. (1975). *Applied multiple regression/correlation analysis for the behavioral sciences.* Oxford, England: Lawrence Erlbaum.

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

Gawronski, B. (2004). Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias. *European Review of Social Psychology*, *15*, 183−217.

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, *132*, 692−731.

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

Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, *54*, 733−740.

Goren, A., & Todorov, A. (2009). Two faces are better than one: Eliminating false trait associations with faces. *Social Cognition*, *27*, 222−248.

Graesser, A. C., Lang, K. L., & Roberts, R. M. (1991). Question answering in the context of stories. *Journal of Experimental Psychology: General*, *120*, 254−277.

Ham, J., & Vonk, R. (2003). Smart and easy: Co-occurring activation of spontaneous trait inferences and spontaneous situational inferences. *Journal of Experimental Social Psychology*, *39*, 434−447.

Hassin, R. R., Aarts, H., & Ferguson, M. J. (2005). Automatic goal inferences. *Journal of Experimental Social Psychology*, *41*, 129−140.

Heider, F. (1958). *The psychology of interpersonal relations.* New York: Wiley.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins, & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133−168). New York: The Guilford Press.

Kammrath, L. K., Mendoza-Denton, R., & Mischel, W. (2005). Incorporating if…then… signatures in person perception: Beyond the person-situation dichotomy. *Journal of Personality and Social Psychology*, *88*, 605−618.

Kelley, H. H. (1973). The process of causal attribution. *American Psychologist*, *28*, 107−128.

Keenan, J. M., Potts, G. R., Golding, J. M., & Jennings, T. M. (1990). Which elaborative inferences are drawn during reading? A question of methodologies. In D. A. Balota, G. B. Flores d'Arcais, & K. Rayner (Eds.), *Comprehension processes in reading* (pp. 377−402). Hillsdale, NJ: Lawrence Erlbaum Associates.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, *95*, 163−182.

Knowles, E. D., Morris, M. W., Chiu, C. Y., & Hong, Y. Y. (2001). Culture and the process of person perception: Evidence for automaticity among East Asians in correcting for situational influences on behavior. *Personality and Social Psychology Bulletin*, *27*, 1344−1356.

Krull, D. S. (1993). Does the grist change the mill? The effect of the perceiver's inferential goal on the process of social inference. *Personality and Social Psychology Bulletin*, *19*, 340−348.

Krull, D. S., & Erickson, D. J. (1995a). Inferential hopscotch: How people draw social inferences from behavior. *Current Directions in Psychological Science*, *4*, 35−38.

Krull, D. S., & Erickson, D. J. (1995b). Judging situations: On the effortful process of taking dispositional information into account. *Social Cognition*, *13*, 417−438.

Lee, F., & Hallahan, M. (2001). Do situational expectations produce situational inferences? The role of future expectations in directing inferences goals. *Journal of Personality and Social Psychology*, *80*, 545−556.

Lieberman, M. D., Gaunt, R., Gilbert, D. T., & Trope, Y. (2002). Reflection and reflexion: A social cognitive neuroscience approach to attributional inference. In M. P. Zanna (Ed.), *Advances in experimental social psychology*, *Vol. 34*. (pp. 199−249)San Diego, CA: Academic Press.

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

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

Malle, B. F. (1999). How people explain behavior: A new theoretical framework. *Personality and Social Psychology Review*, *3*, 23−48.

Malle, B. F. (2008). Is there a hierarchy of social inferences? The ease and speed of judging intentionality, mind, and personality. Paper presented at the annual meeting of the Person Memory Interest Group, Petaluma, CA.

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 82−91.

Molden, D. C., & Dweck, C. S. (2006). Finding "meaning" in psychology: A lay theories approach to self-regulation, social perception, and social development. *American Psychologist*, *61*, 192−203.

Molden, D. C., Plaks, J. E., & Dweck, C. S. (2006). "Meaningful" social inferences: Lay theories and inferential processes. *Journal of Experimental Social Psychology*, *42*, 738−752.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, *132*, 297−326.

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, *114*, 510−532.

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods*, 2nd Ed. Thousand Oaks, CA: Sage.

Read, S. J., & Miller, L. C. (1993). Rapist or "regular guy": Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin*, *19*, 526−541.

Read, S. J., & Miller, L. C. (1998). On the dynamic construction of meaning: An interactive activation and competition model of social perception. In S. J. Read, & L. C. Miller (Eds.), *Connectionist models of social reasoning and social behavior* (pp. 27−68). Mahwah, NJ: Erlbaum.

Reeder, G. D. (2009). Mindreading: Judgments about intentionality and motives in dispositional inference. *Psychological Inquiry*, *20*, 1−18.

Reeder, G. D., Vonk, R., Ronk, M. J., Ham, J., & Lawrence, M. (2004). Dispositional attribution: Multiple inferences about motive-related traits. *Journal of Personality and Social Psychology*, *86*, 530−544.

Ross, L. D., & Nisbett, R. E. (1991). *The person and the situation: Perspectives of social psychology.* New York: McGraw-Hill.

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.

Swinney, D. A. (1979). Lexical access during comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, *18*, 645−659.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, *12*, 435−467.

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

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

Todorov, A., & Uleman, J. S. (2004). The person reference process in spontaneous trait inferences. *Journal of Personality and Social Psychology*, *87*, 482−493.

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

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

Trope, Y., & Gaunt, R. (2000). Processing alternative explanations of behavior: Correction or integration? *Journal of Personality and Social Psychology*, *79*, 344−354.

Trope, Y., & Gaunt, R. (2003). Attribution and person perception. In M. A. Hogg, & J. Cooper (Eds.), *The sage handbook of social psychology* (pp. 190−209). New York: Sage Publications.

Uleman, J. S. (1988). [Trait and gist inference norms for over 300 potential trait-implying sentences]. New York University, Unpublished data.

Uleman, J. S. (1999). Spontaneous versus intentional inferences in impression formation. In S. Chaiken, & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 141−160). New York: Guilford Press.

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

Uleman, J. S., Hon, A., Roman, R., & Moskowitz, G. B. (1996). On-line evidence for spontaneous trait inferences at encoding. *Personality and Social Psychology Bulletin*, *22*, 377−394.

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

Uleman, J. S., Saribay, S. A., & Gonzalez, C. M. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual Review of Psychology*, *59*, 329−360.

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

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