ARTICLE IN PRESS

Journal of Experimental Social Psychology xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

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

Does contextualized attitude change depend on individual differences in responses to belief-incongruent information?☆

Skylar M. Brannon*, Bertram Gawronski

University of Texas at Austin, United States

ARTICLE INFO

Handling editor: Elizabeth Page-Gould

Keywords:
Attitude change
Context effects
Implicit theories of personality
Need for structure
Preference for consistency

ABSTRACT

Previous research has shown that changes in the evaluation of an attitude object can be limited to the context in which counterattitudinal information was learned. To account for these findings, it has been proposed that exposure to expectancy-violating information enhances attention to context, which leads to an integration of the context into the representation of expectancy-violating counterattitudinal information. Although a considerable body of evidence supports these assumptions, it is still unclear whether contextualized attitude change is a general phenomenon that is robust across individuals or instead depends on psychological characteristics of the perceiver. To address this question, the current research tested whether contextualized attitude change is moderated by three individual difference variables that are known to influence responses to belief-incongruent information: preference for consistency, need for structure, and implicit theories of personality. Based on the hypothesis that contextualized attitude change is due to enhanced attention to context during encoding of expectancy-violating information, we hypothesized that individual differences along the three dimensions should moderate contextualized attitude change via differences in attention to context in response to expectancy-violating information. Contrary to this hypothesis, none of the three variables moderated contextualized attitude change (Experiments 1 and 2) and attention to context during exposure to expectancy-violating information (Experiment 3). Implications for the generality of contextualized attitude change, research on the three individual difference variables, and cognitive consistency more broadly are discussed.

1. Introduction

People often behave inconsistently. One might observe a new colleague being nice and friendly at work, but later find the same colleague being nasty and rude at a grocery store. How do observers account for such inconsistencies in their evaluations of others? Previous research suggests that changes in the evaluation of another person in response to counterattitudinal information about that person can be limited to the context in which the counterattitudinal information was learned (for reviews, see Gawronski & Cesario, 2013; Gawronski et al., 2018). That is, evaluations may reflect newly learned counterattitudinal information only in the context in which this information was learned and the valence of initial attitudinal information in any other context. To explain such patterns of contextualized attitude change, it has been proposed that expectancy-violating information enhances attention to context, which leads to an integration of the context into the representation of expectancy-violating counterattitudinal information (Gawronski, Rydell, Vervliet, & De Houwer, 2010). Although these assumptions are supported by a considerable body of evidence (for a review, see Gawronski et al., 2018), it is still unclear whether contextualized attitude change is a general phenomenon that is robust across individuals or instead depends on psychological characteristics of the perceiver.

Drawing on the hypothesis that contextualized attitude change is due to enhanced attention to context during the encoding of expectancy-violating information (Gawronski et al., 2010), the current research tested whether contextualized attitude change is moderated by three individual difference variables that are known to influence responses to belief-incongruent information: preference for consistency (PFC; Cialdini, Trost, & Newsom, 1995), need for structure (NFS; Neuberg & Newsom, 1993), and implicit theories of personality (ITP; Chiu, Hong, & Dweck, 1997). Across three experiments, we tested whether individual differences along the three dimensions moderate contextualized changes in evaluations of another person (Experiments 1

☆ The research reported in this article has been partly supported by grants from the Canada Research Chairs Program (215983) and the Social Sciences and Humanities Research Council of Canada (410-2011-0222). We would like to thank Allison DeJong, Nicole Dryburgh, Stephanie Gonzalez, Sainsanaa Khurelbaatar, Grace Ki, Nicole Miles, Nick Stiles, Zachary Taylor, and Ting Zhao for their assistance in collecting the data.

* Corresponding author.
E-mail address: skylar_brannon@utexas.edu (S.M. Brannon).

https://doi.org/10.1016/j.jesp.2018.03.015
Received 18 September 2017; Received in revised form 30 March 2018; Accepted 30 March 2018

and 2) and attention to context during exposure to expectancy-violating information (Experiment 3).

1.1. Contextualized attitude change

To reconcile mixed findings regarding the malleability of attitudes, Gawronski et al. (2018) suggested that whether or not attitudes appear resistant to counterattitudinal information can depend on the context in which evaluations are measured. Specifically, they suggested that evaluations of an attitude object might reflect newly learned counterattitudinal information only when measured in the context in which the counterattitudinal information was learned. Yet, evaluations may continue to reflect initially learned attitudinal information when measured in the context in which the initial attitudinal information was learned or a novel context in which the attitude object has not been encountered before (for a review of similar findings in research on animal learning, see Bouton, 2004). For example, if Megan observes her new colleague Don being nice and friendly at work, and later observes Don being nasty and rude at a grocery store, her evaluation of Don may reflect the new, negative information only within the context of the grocery store. Conversely, the initial, positive information may continue to influence her evaluation of Don within the work context as well as any novel context in which she has not encountered him before (e.g., a resort).

Rydell and Gawronski (2009) provided the first evidence that social attitudes show such patterns of contextualized attitude change (for a meta-analysis, see Gawronski, Hu, Rydell, Vervliet, & De Houwer, 2015). In a series of studies, participants first formed an impression of a target individual based on statements describing either positive or negative behaviors. The statements were paired with a picture of the target, both of which were presented against a colored background (e.g., blue). Subsequently, participants learned new information about the target that was evaluatively incongruent with the initial information. The new statements were paired with the same picture against a different colored background (e.g., yellow). Finally, participants completed an affective priming task (Payne, Cheng, Govorun, & Stewart, 2005) to measure spontaneous evaluations of the target. Critically, the target's picture was presented against three different colored backgrounds to assess whether participants' evaluations differed across contexts: the background of the initial attitudinal information (e.g., blue), the background of the counterattitudinal information (e.g., yellow), and a novel background that was not part of the impression formation task (e.g., green). Rydell and Gawronski (2009) found that evaluations reflected the counterattitudinal information only when the target was presented against the background in which the counterattitudinal information was learned. In contrast, evaluations reflected the initial attitudinal information when the target was presented against the background in which the initial attitudinal information was learned. Moreover, when the target was presented against a novel background that was not part of the impression formation task, evaluations again reflected the initial attitudinal information.

1.2. Representational theory

To account for context-dependent changes in evaluations, Gawronski et al. (2010) proposed a representational theory that explains why the effect of counterattitudinal information is sometimes limited to the context in which this information was learned (for a review, see Gawronski et al., 2018). A central assumption of this theory is that attention to context during the learning of evaluative information determines whether the context is integrated into the mental representation of that information. Gawronski et al. (2010) further suggested that attention to context is typically low when encoding initial attitudinal information about an object (see Gilbert & Malone, 1995), leading initial attitudinal information to be stored in a context-free representation. Moreover, because attention to context is typically enhanced by exposure to expectancy-violating information (see Roese & Sherman, 2007), expectancy-violating counterattitudinal information is assumed to be stored in a contextualized representation. Thus, after learning attitude-incongruent information about an object, the mental representation of that object takes on a "dual" nature by including (1) a context-free representation of the initial attitudinal information and (2) contextualized representation of the counterattitudinal information. As a result, evaluative responses to the object should differ depending on the presence versus absence of the context in which the counterattitudinal information was learned. In line with the principle of pattern matching in memory activation (Smith, 1996), the contextualized representation of counterattitudinal information should be activated when the object is encountered in the context in which the counterattitudinal information was learned. Conversely, the context-free representation of initial attitudinal information should be activated when the object is encountered in a context that is different from the context in which the counterattitudinal information was learned.

In addition to providing an explanation of contextualized attitude change, Gawronski et al.'s (2010) representational theory also includes specific predictions about the conditions under which contextualized attitude change should *not* occur. Specifically, the theory suggests that contextualized attitude change should be eliminated when attention to context is low during both the encoding of initial attitudinal information and the encoding of counterattitudinal information. In this case, the two kinds of information should be integrated in a single context-free representation, leading to evaluative responses that reflect a mixture of attitudinal and counterattitudinal information regardless of the context. Consistent with the proposed role of attention to context, Gawronski, Ye, Rydell, and De Houwer (2014) demonstrated that attention to incidental context cues (i.e., background color of a computer screen) is relatively high during the encoding of attitude-incongruent information, but relatively low during the encoding of attitude-congruent information (see also Brannon & Gawronski, in press; Brannon, Sacchi, & Gawronski, 2017; Ye, Tong, Chiu, & Gawronski, 2017). Moreover, Gawronski et al. (2010) found that contextualized attitude change was fully eliminated when attention to context during the encoding of attitude-incongruent information had been experimentally reduced. In this case, evaluations reflected an equally weighted mixture of attitudinal and counterattitudinal information regardless of the context.

Although these and various other findings support the assumptions of the representational theory (for a review, see Gawronski et al., 2018), it is still unclear whether contextualized attitude change is a general phenomenon that is robust across individuals or instead depends on psychological characteristics of the perceiver. Drawing on the hypothesis that contextualized attitude change is due to enhanced attention to context during the encoding of expectancy-violating information, the current research tested whether contextualized attitude change is moderated by three individual difference variables that are known to influence responses to belief-incongruent information.

1.3. Individual differences in responses to inconsistency

According to the representational theory, a critical factor underlying contextualized attitude change is mental conflict in response to belief-incongruent information, which is known to elicit a broad range of cognitive, affective, and motivational reactions (Festinger, 1957; Gawronski & Brannon, in press; Proulx, Inzlicht, & Harmon-Jones, 2012). These reactions are assumed to be central for contextualized attitude change, in that they involve enhanced attention to context during the encoding of expectancy-violating counterattitudinal information. Classic theories proposed that inconsistency between cognitive elements is inherently aversive (e.g., Festinger, 1957), which motivates people to reconcile the inconsistency. However, later work suggested that people differ in their tolerance for inconsistency and, thus, the extent to which they show cognitive, affective, and motivational reactions in response to inconsistent information. This work has largely focused on three trait dimensions that have been shown to moderate responses to inconsistency: PFC (Cialdini et al., 1995), NFS (Neuberg & Newsom, 1993), and ITP (Chiu et al., 1997). Based on the processes suggested by Gawronski et al.'s (2010) representational theory, these personality dimensions may also influence the emergence of contextualized attitude change.

1.3.1. PFC

Cialdini et al. (1995) argued that people differ in the extent to which they prefer consistency in their own and others' behaviors. They further suggested that these differences could account for disparate findings regarding responses to inconsistency. In line with these assumptions, they found that individuals high in PFC showed behaviors aimed toward reconciling an inconsistency (e.g., attitude change after freely choosing to write a counter-attitudinal message; see Festinger & Carlsmith, 1959). Conversely, individuals low in PFC did not show any such behaviors. Moreover, individuals high in PFC have been found to experience more negative arousal in response to attitudinal ambivalence than those low in PFC (Newby-Clark, McGregor, & Zanna, 2002). Together, these findings suggest that individuals low in PFC may show less extreme reactions to expectancy-violating information compared to individuals high in PFC. As a result, individuals low in PFC may pay less attention to context following an expectancy-violation than individuals high in PFC, which should lead to attenuated patterns of contextualized attitude change among individuals low in PFC.

1.3.2. NFS

People also differ in the extent to which they want to avoid ambiguity and organize their surroundings into structured, manageable chunks of information, a personality dimension known as NFS (Neuberg & Newsom, 1993). High NFS leads people to organize incoming information in terms of broad strokes, creating schemas or stereotypes that meaningfully categorize disparate pieces of information into more integrative structures of knowledge. Although NFS and PFC are conceptually distinct constructs, they are similar in the sense that both reflect an individual's need for stable and coherent knowledge structures (Kruglanski, 1990; Kruglanski & Webster, 1996). Thus, like PFC, individual differences in NFS are associated with differences in the processing of inconsistency, in that individuals low in NFS are more tolerant of inconsistencies than individuals high in NFS. For example, individuals low in NFS respond to expectancy-violations with increased cognitive flexibility and creativity, whereas those high in NFS become more cognitively rigid in response to expectancy-violations (Gocłowska, Baas, Crisp, & De Dreu, 2014). Based on these findings, one could argue that individuals low in NFS may pay less attention to context following an expectancy-violation than individuals high in NFS, which should lead to attenuated patterns of contextualized attitude change among individuals low in NFS.

1.3.3. ITP

Finally, individuals vary in the degree to which they believe people's personality traits are malleable (Dweck & Leggett, 1988), and these beliefs have been shown to influence the processing of social information (for a review, see Plaks, 2017). In contrast to the need for stable and coherent knowledge structures associated with PFC and NFS, ITP refers more specifically to beliefs about personality traits and their malleability. Nevertheless, the three constructs converge in terms of their relation to one's preference and expectation that others behave in coherent and predictable ways. For example, people who believe a person's traits are relatively fixed (i.e., entity theorists) display greater attention to and better memory for stereotype-consistent information. Conversely, those who believe a person's traits are relatively malleable (i.e., incremental theorists) display greater attention to and better memory for stereotype-inconsistent information (Plaks, Grant, & Dweck, 2005; Plaks, Stroessner, Dweck, & Sherman, 2001). Applied to the current question, these findings suggest that incremental theorists may show attenuated patterns of contextualized attitude change compared to entity theorists. For incremental theorists, observing inconsistency in a person's behavior should not be particularly surprising, given that people's personality is assumed to be malleable. In contrast, entity theorists should be surprised by observing inconsistency in a person's behavior, given that people's personality is assumed to be fixed. Hence, incremental theorists may pay less attention to context following an expectancy-violation than entity theorists, which should lead to attenuated patterns of contextualized attitude change among individuals holding an incremental theory.

1.4. The current research

Although previous research on PFC, NFS, and ITP has not specifically investigated their link with attention to context, such a link can be inferred from extant research regarding the relation between the three constructs and responses to inconsistency. To the extent that individuals high in PFC, individuals high in NFS, and individuals who hold an entity theory show more extreme reactions to inconsistency, individuals with these personality traits might also pay more attention to context in response to expectancy-violating counterattitudinal information. Conversely, to the extent that individuals low in PFC, individuals low in NFS, and individuals who hold an incremental theory have a greater tolerance for inconsistency, attention to context in response to expectancy-violating counterattitudinal information should be relatively low for individuals with these personality traits. As a result, contextualized attitude change should be more pronounced among individuals high in PFC, individuals high in NFS, and individuals holding an entity theory of personality. In contrast, contextualized attitude change should be less pronounced (or eliminated) among individuals low in PFC, individuals low in NFS, and individuals holding an incremental theory of personality. In the current research, we tested these predictions by investigating whether contextualized attitude change is moderated by individual differences in PFC, NFS, and ITP (Experiments 1 and 2). In a follow-up study, we also tested the hypothesis that individual differences along the three dimensions influence contextualized attitude change by moderating the amount of attention to context following an expectancy-violation (Experiment 3).

Counter to our predictions, none of the three individual difference variables moderated contextualized attitude change and attention to context following an expectancy-violation. Although we replicated the typical pattern of contextualized attitude change, individual differences in PFC, NFS, and ITP did not moderate contextualized attitude change. The same was true for attention to context following an expectancy-violation. Although we replicated earlier findings showing high attention to context during the encoding of attitude-incongruent, but not attitude-congruent, information (see Brannon et al., 2017; Brannon & Gawronski, in press; Gawronski et al., 2014; Ye et al., 2017), attention to context following an expectancy-violation was unrelated to the three individual difference variables. These findings have important implications for the generality of contextualized attitude change, research on the three individual difference constructs, and cognitive consistency more broadly, which we will explain in the General Discussion section.[1]

2. Experiment 1

Experiment 1 adapted Rydell and Gawronski's (2009) impression formation paradigm to investigate whether individual differences in PFC, NFS, and ITP moderate contextualized attitude change. Participants first completed measures of PFC, NFS, and ITP. Subsequently,

[1] For all studies reported here, the data were collected in one shot without prior statistical analyses. We report all data exclusions, all measures, and all manipulations. All materials, data, and analysis files are available at https://osf.io/usgz3/.

they formed impressions of two individuals, one being portrayed as negative and one being portrayed as positive. The initial information about these individuals was displayed against the same colored background. In a second block of the impression formation task, participants were presented with new information about each individual that was evaluatively incongruent with the information provided in the first block. This counterattitudinal information was presented against a colored background that differed from the one against which the initial information was presented. To measure evaluations of the two individuals, participants completed a speeded evaluation task (Ranganath, Smith, & Nosek, 2008) in which they were shown pictures of each individual and asked to indicate as quickly as possible whether they felt positively or negatively about the presented individual. Critically, each picture was presented against each of three different backgrounds: the background against which the initial attitudinal information was displayed (first learning context), the background against which the counterattitudinal information was displayed (second learning context), and a novel background that did not appear during the impression formation task (novel context). These evaluations served as the primary dependent variable and were analyzed as a function of target individual and measurement context (see Gawronski et al., 2014). In this design, evidence for contextualized attitude change is reflected in an interaction pattern involving (1) a stronger impact of the counterattitudinal information when evaluations are measured in the second learning context compared to evaluations measured in the first learning context and (2) a stronger impact of the counterattitudinal information when evaluations are measured in the second learning context compared to evaluations measured in a novel context (see Fig. 1). Our main question was whether the predicted pattern of contextualized attitude change depends on participants' level of PFC, NFS, and ITP, respectively.

2.1. Method

2.1.1. Participants and design

Ninety-one psychology undergraduates completed the study in exchange for course credit. Due to experimenter error, data were lost for five participants, resulting in a final sample of 86 participants (57 women, 29 men; $M_{\text{age}} = 18.40$ years, $SD_{\text{age}} = 0.73$ years).[2] Participants were randomly assigned to the four conditions of a 2 (Color Order: yellow-blue vs. blue-yellow) × 2 (Target-Valence Matching in First Block: target1-positive, target2-negative vs. target2-positive, target1-negative) between-subjects design.

2.1.2. Individual difference measures

Participants first completed individual difference measures of PFC, NFS, and ITP.

2.1.2.1. PFC. Participants first indicated their agreement with the 18 statements of the PFC Scale (Cialdini et al., 1995). Sample statements include "I want my close friends to be predictable" and "I typically prefer to do things the same way." Participants rated their agreement on 9-point scales ranging from 1 (*Strongly Disagree*) to 9 (*Strongly Agree*). Responses were aggregated such that higher scores reflect higher PFC (Cronbach's $\alpha = 0.87$).

2.1.2.2. NFS. Then, participants indicated their agreement with the 11 statements comprising the NFS Scale (Neuberg & Newsom, 1993). Sample statements include "It upsets me to go into a situation without knowing what I can expect from it" and "I hate to be with people who are unpredictable." Participants rated their agreement on 6-point scales with the response options 1 (*Strongly Disagree*), 2 (*Disagree*), 3 (*Somewhat Disagree*), 4 (*Somewhat Agree*), 5 (*Agree*), 6 (*Strongly Agree*). Responses were aggregated such that higher scores reflect a higher NFS (Cronbach's $\alpha = 0.82$).

2.1.2.3. ITP. Finally, participants completed eight items to measure ITP (Chiu et al., 1997). Sample items include "Everyone is a certain kind of person and there is not much that can be done to really change that" and "The kind of person someone is something very basic about them and it can't be changed very much." Participants rated their agreement with eight statements on the 6-point scales with the response options 1 (*Strongly Disagree*), 2 (*Disagree*), 3 (*Somewhat Disagree*), 4 (*Somewhat Agree*), 5 (*Agree*), 6 (*Strongly Agree*). Responses were aggregated such that higher scores reflect stronger endorsement of an entity theory and lower scores reflect stronger endorsement of an incremental theory (Cronbach's $\alpha = 0.87$).

2.1.3. Impression formation task

Participants were asked to form impressions of two individuals based on statements describing their behavior. The statements were adapted from Rydell and Gawronski (2009). Each statement was presented below a picture of one of the two individuals for 5000 ms. In a first learning block, one of the two individuals was paired with 25 positive statements, the other target was paired with 25 negative statements. Which individual was described as positive or negative was counter-balanced across participants. For half of the participants, the 50 statements were presented against a yellow background. For the other half, the 50 statements were presented against a blue background. In a second block of the impression formation task, participants were told that they would now receive additional information about each individual. In this block, the valence of the information about each individual was reversed. That is, the individual that was initially paired with 25 positive statements was now presented with 25 negative statements, and vice versa. Additionally, participants who viewed the initial statements against a yellow background viewed the new statements against a blue background, and vice versa. As in the first block of the impression formation task, each statement was presented below a picture of one of the target individuals for 5000 ms.

2.1.4. Evaluation measure

After the impression formation task, participants completed a speeded evaluation task (Ranganath et al., 2008) to measure their spontaneous evaluations of the two target individuals. Participants were told that they would be briefly presented with faces and that they should indicate their immediate "gut responses" toward each stimulus. Each trial began with a blank screen for 500 ms, followed by a fixation cross for 500 ms, and, finally, one of the faces from the impression formation task for 100 ms. The presentation of each face was followed by a blank screen for 100 ms and a prompt for a response for 800 ms. Participants had 1000 ms (beginning with the onset of the stimulus) to indicate whether they felt either positively or negatively about the individual by pressing a right-hand key (*Numpad 5*) for positive responses and a left-hand key (*A*) for negative responses. If participants did not respond by the end of the response window, a message appeared on the screen asking them to respond faster. To assess whether evaluations of the two individuals depended on context, each face was presented against the background of the first learning block (i.e., blue or yellow), the background of the second learning block (i.e., yellow or blue), or a novel background that was not part of the impression formation task (i.e., green). Each face was presented ten times against each of the three background colors, summing up to a total of 60 trials.

[2] Our initial analysis plan involved calculating two scores reflecting the difference in evaluations of a given target in (1) the second learning context vs. the first learning context and (2) the second learning context vs. a novel context (see Gawronski et al., 2015). To test moderation of contextualized attitude change by individual differences in PFC, NFC, and ITP, we planned to correlate the two difference scores with the three individual difference measures. Our power calculations for Experiments 1 and 2 were based on these planned analyses, which, in response to feedback by the editor, have been replaced by analyses using linear mixed effect models. A sample size of 86 provides 80% power to detect a correlation of $r = 0.29$ between each difference score and the three individual difference measures.

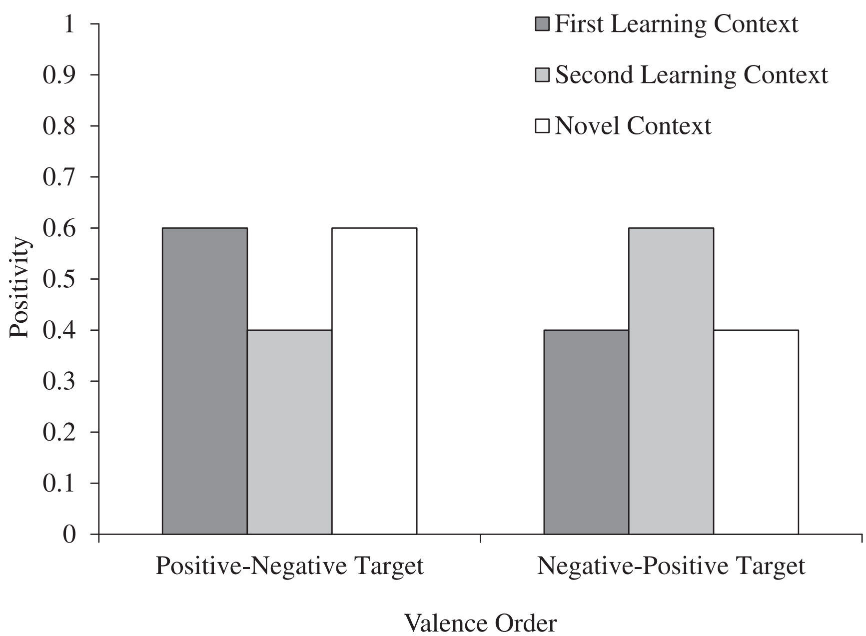

Fig. 1. Hypothetical pattern of target evaluations as a function of valence order (positive-negative vs. negative-positive) and measurement context (first learning context vs. second learning context vs. novel context). Contextualized attitude change is reflected in (1) a stronger impact of the second (counterattitudinal) information when evaluations are measured in the second learning context compared to evaluations measured in the first learning context and (2) a stronger impact of the second (counterattitudinal) information when evaluations are measured in the second learning context compared to evaluations measured in a novel context.

2.2. Results

Evaluation scores were aggregated by first eliminating trials on which participants failed to respond within the 1000 ms response window and then calculating the proportion of valid trials on which participants made a *positive* judgment for each target individual against each of the three background colors. Thus, higher scores reflect more favorable evaluations of a given target within each of the three contexts.

To test whether contextualized attitude change is moderated by individual differences in NFS, PFC, and ITP, we submitted participants' aggregated evaluation scores to separate linear mixed effects (LME) models for each of the three individual difference measures.[3] Along with the respective individual difference measure, Valence Order (positive vs. negative) and Measurement Context (first learning context vs. second learning context vs. novel context) were entered as predictors of participants' evaluation scores. In line with past research, we expected to obtain a significant two-way interaction between Valence Order and Measurement Context, reflecting the predicted pattern of contextualized attitude change (see Fig. 1). Further, we expected this interaction to be qualified by a higher-order interaction with each of the three individual difference measures, indicating that the emergence of contextualized attitude change depends on each of the individual difference measures. Specifically, we expected contextualized attitude change to be attenuated for participants low in PFC, participants low in NFS, and participants with an incremental theory of personality. Conversely, we expected contextualized attitude change to be enhanced for participants high in PFC, participants high in NFS, and participants with an entity theory of personality.

2.2.1. PFC

First, evaluation scores were submitted to a Valence Order × Measurement Context × PFC model, with Measurement Context as a within-subjects factor, Valence Order as a between-subjects factor, and PFC as a continuous predictor (see Table 1). Contrary to our hypotheses, the three-way interaction was not significant and was removed from the model for a more accurate estimation of lower-order interactions and main effects (see Kutner, Nachtsheim, Neter, & Li, 2005). In the reduced model, neither the two-way interaction between Valence Order and PFC, nor the two-way interaction between Measurement Context and PFC were significant. These interactions were, thus, removed from the model, and the model was re-estimated. The final model revealed in a significant two-way interaction between Valence Order and Measurement Context. To decompose this interaction, a priori pairwise contrasts were conducted.[4] These contrasts suggest a significant pattern of contextualized attitude change in both Valence Order conditions (see Fig. 2). For positive-negative targets, evaluations were more positive in the first learning context as compared to the second learning context, χ^2 (1) = 6.70, p = .010, r = 0.28. Moreover, positive-negative targets were evaluated more positively in the novel context than in the second learning context, χ^2 (1) = 5.41, p = .020, r = 0.25. Conversely, negative-positive targets were evaluated more negatively in the first learning context as compared to the second learning context, $\chi^2(1)$ = 14.53, p < .001, r = 0.41. Moreover, negative-positive targets were evaluated more negatively in the novel context than in the second learning context, χ^2 (1) = 4.32, p = .038, r = 0.22.

2.2.2. NFS

We repeated the full LME model with NFS (see Table 1). As for PFC, the three-way interaction between Valence Order, Measurement Context, and NFS was not significant and, thus, was removed from the model. In the reduced model, the two-way interaction between Measurement Context and NFS was not significant and, thus, was removed from the model. The final model revealed in a significant main effect of Valence Order, which was qualified by higher-order interactions with both NFS and Measurement Context. Because the two-way interaction between Valence Order and NFS is independent of Measurement Context and, thus, irrelevant for our predictions regarding contextualized attitude change, we do not discuss it further. The breakdown of the

[3] Correlation analyses with the three individual difference measures revealed a significant positive correlation between PFC and NFS (r = 0.54, p < .001). There were no significant relations between PFC and ITP (r = 0.12, p = .286) and between NFS and ITP (r = 0.15, p = .172).

[4] Decomposition of the interaction between Valence Order and Measurement Context in Experiments 1 and 2 were conducted using the phia package (De Rosario-Martinez, 2015) in RStudio (RStudio Team, 2016). Effect sizes for these contrasts were calculated using the online companion for Lipsey and Wilson's (2001) guide to meta-analysis (https://www.campbellcollaboration.org/escalc/html/EffectSizeCalculator-R5.php).

Table 1
LME model results and Bayes factors, Experiment 1.

Model	Predictor	Model *df*	Residual *df*	*F*	*p*	R_p^2	BF_{10}	BF_{10} interpretation
PFC	Valence Order	1	425	0.12	.730	0.000	0.10	Substantial evidence for H_0
	Measurement Context	2	425	0.62	.540	0.003	0.04	Strong evidence for H_0
	PFC	1	84	1.72	.193	0.020	0.21	Substantial evidence for H_0
	Valence Order × Measurement Context	2	425	10.68	< .001	0.048	1077.20	Decisive evidence for H_1
	Valence Order × PFC	1	422	2.27	.132	0.005	0.46	Anecdotal evidence for H_0
	Measurement Context × PFC	2	422	0.09	.917	0.000	0.03	Very strong evidence for H_0
	Valence Order × Measurement Context × PFC	2	420	1.46	.232	0.007	0.26	Substantial evidence for H_0
NFS	Valence Order	1	424	5.33	.021	0.012	0.10	Substantial evidence for H_0
	Measurement Context	2	424	0.62	.538	0.003	0.04	Strong evidence for H_0
	NFS	1	84	0.34	.561	0.004	0.13	Substantial evidence for H_0
	Valence Order × Measurement Context	2	424	10.75	< .001	0.048	1059.29	Decisive evidence for H_1
	Valence Order × NFS	1	424	5.78	.017	0.013	2.53	Anecdotal evidence for H_1
	Measurement Context × NFS	2	422	0.14	.869	0.001	0.03	Very strong evidence for H_0
	Valence Order × Measurement Context × NFS	2	420	0.60	.548	0.003	0.11	Substantial evidence for H_0
ITP	Valence Order	1	425	0.12	.729	0.000	0.10	Substantial evidence for H_0
	Measurement Context	2	425	0.62	.538	0.003	0.04	Strong evidence for H_0
	ITP	1	84	3.91	.051	0.044	0.42	Anecdotal evidence for H_0
	Valence Order × Measurement Context	2	425	10.73	< .001	0.048	1100.45	Decisive evidence for H_1
	Valence Order × ITP	1	422	2.07	.151	0.005	0.40	Anecdotal evidence for H_0
	Measurement Context × ITP	2	422	0.72	.488	0.003	0.05	Strong evidence for H_0
	Valence Order × Measurement Context × ITP	2	420	0.12	.889	0.001	0.07	Strong evidence for H_0

Note. LME analyses were conducted using the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2016) in RStudio (RStudio Team, 2016) using Kenward-Roger approximated degrees of freedom. Residual degrees of freedom for each effect reflect removal of non-significant higher order effects, which were removed from the model for a more accurate estimation of lower order interactions and main effects (e.g., Kutner et al., 2005). Bayes factors were obtained using the BayesFactor package (Morey & Rouder, 2015) in RStudio (RStudio Team, 2016). To obtain the Bayes factor, a model including only lower level effects (or an empty model for the main effects) was compared to a model containing both the lower level effects and the effect of interest. Evidence category labels for Bayes Factors follow recommendations from Wetzles and Wagenmakers (2012).

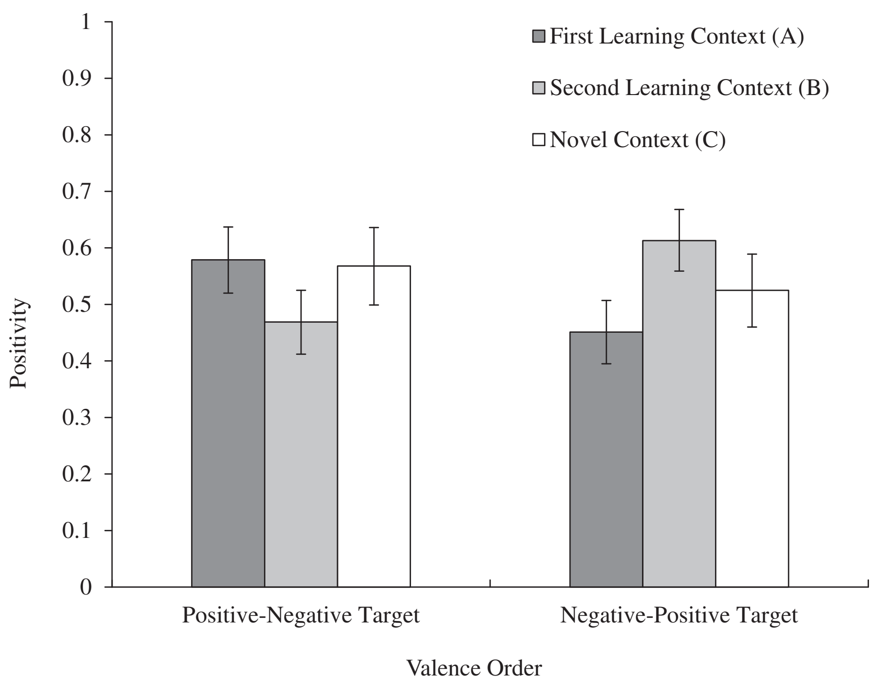

Fig. 2. Target evaluations as a function Valence Order (positive-negative vs. negative-positive) and Measurement Context, Experiment 1. Higher scores reflect more positive evaluations. Error bars represent 95% confidence intervals. Values represent descriptives at the group level collapsed across levels of PFC, NFS, and ITP.

Valence Order by Measurement Context interaction was consistent with the one in the model containing PFC (see Fig. 2). For positive-negative targets, evaluations were more positive in the first learning context as compared to the second learning context, χ^2 (1) = 6.71, p = .010, r = 0.28. Moreover, positive-negative targets were evaluated more positively in the novel context than in the second learning context, χ^2 (1) = 5.42, p = .020, r = 0.25. Conversely, negative-positive targets were evaluated more negatively in the first learning context as compared to the second learning context, χ^2 (1) = 14.56, p < .001, r = 0.41. Moreover, negative-positive targets were evaluated more negatively in the novel context than in the second learning context, χ^2 (1) = 4.33, p = .037, r = 0.22.

2.2.3. ITP

Finally, the LME model was repeated for ITP (see Table 1). As with the other two models, the three-way interaction between Valence Order, Measurement Context, and ITP was not significant. This interaction was removed from the model, and the model was re-estimated.

In the reduced model, neither the two-way interaction between Valence Order and ITP nor the two-way interaction between Measurement Context and ITP was significant. These interactions were, thus, removed from the model, and the model was re-estimated. The final model revealed a significant main effect of ITP that is irrelevant for our predictions regarding contextualized attitude change. More important, the interaction between Valence Order and Measurement Context was significant. As in the previous models, this interaction suggested a significant pattern of contextualized attitude change in both Valence Order conditions (see Fig. 2). For positive-negative targets, evaluations were more positive in the first learning context as compared to the second learning context, χ^2 (1) = 6.67, p = .010, r = 0.28. Moreover, positive-negative targets were evaluated more positively in the novel context than in the second learning context, χ^2 (1) = 5.38, p = .020, r = 0.26. Conversely, negative-positive targets were evaluated more negatively in the first learning context as compared to the second learning context, χ^2 (1) = 14.47, $p < .001$, r = 0.41. Moreover, negative-positive targets were evaluated more negatively in the novel context than in the second learning context, χ^2 (1) = 4.30, p = .039, r = 0.22.

2.2.4. Bayesian analysis

The LME analysis replicated earlier evidence for contextualized attitude change, as evidenced by the interaction between Valence Order and Measurement Context (see Fig. 2). Contrary to our hypothesis, however, the two-way interaction between Valence Order and Measurement Context was not qualified by a higher-order interaction with any of the three individual difference measures. This suggests that contextualized attitude change was not moderated by individual differences in PFC, NFS, and ITP. This conclusion, however, is based on null-hypothesis statistical test (NHST) procedures, which do not provide any indication of the extent to which the alternative hypothesis should be favored over the null, and vice versa. To address concerns about inferences from non-significant effects, we conducted Bayesian analyses that allow for more direct conclusions regarding the null hypothesis (see Table 1). To the extent that Bayesian analyses corroborate the results of the NHSTs, they would help to rule out low statistical power as an alternative explanation for the non-significant three-way interactions in the NHSTs. Supporting the LME analyses reported above, there was decisive evidence for contextualized attitude change, as captured by the two-way interaction between Valence Order and Measurement Context. Moreover, there was strong to substantial evidence that the two-way interaction between Valence Order and Measurement Context was not qualified by either PFC, NFS, or ITP.

2.3. Discussion

Experiment 1 investigated whether individual differences in PFC, NFS, and ITP moderate contextualized attitude change. Replicating past research (e.g., Gawronski et al., 2014), evaluations of a target individual reflected newly learned counterattitudinal information only in the context in which the counterattitudinal information was learned. In contrast, evaluations reflected initial attitudinal information in the context in which the attitudinal information was learned and a novel context in which the target had not been encountered before. Contrary to our hypotheses regarding individual differences in responses to inconsistency, these patterns were not moderated by PFC, NFS, or ITP.

3. Experiment 2

Although the results of Experiment 1 suggest that individual differences in PFC, NFS, and ITP do not moderate contextualized attitude change, the sample size was relatively small. To address this limitation, Experiment 2 aimed to replicate the findings of Experiment 1 with a larger sample size using the same materials.

3.1. Method

3.1.1. Participants and design

One hundred eighty-four psychology undergraduates (108 women, 75 men; M_{age} = 18.94 years, SD_{age} = 1.21 years; demographic data missing for one participant) completed the study in exchange for course credit.[5] As in Experiment 1, participants were randomly assigned to the four conditions of a 2 (Color Order) × 2 (Target-Valence Matching in First Block) between-subjects design.

3.1.2. Procedures and measures

All aspects of Experiment 2 were identical to those of Experiment 1. The measures of PFC, NFS, and ITP showed satisfactory reliabilities with Cronbach's αs of 0.88, 0.82, and 0.90, respectively.

3.2. Results

As in Experiment 1, we submitted participants' aggregated evaluation scores to separate LME models for each individual difference measure.[6]

3.2.1. PFC

First, we conducted an LME analysis using the full Valence Order × Measurement Context × PFC model, with Measurement Context as within-subjects factor, Valence Order as a between-subjects factor, and PFC as a continuous predictor (see Table 2). The three-way interaction was not significant and, thus, was removed from the model. In the reduced model, neither the two-way interaction between Valence Order and PFC nor the two-way interaction between Measurement Context and PFC were significant. These two interactions were, thus, removed from the model, and the model was re-estimated. The final model resulted in a significant main effect of Valence Order, which was qualified by a significant two-way interaction between Valence Order and Measurement Context. To decompose this interaction, we again conducted a priori pairwise contrasts. Replicating the findings of Experiment 1, these contrasts suggested a pattern of contextualized attitude change in both Valence Order conditions (see Fig. 3). For positive-negative targets, evaluations were more positive in the first learning context as compared to the second learning context, χ^2 (1) = 4.04, p = .044, r = 0.15. Moreover, positive-negative targets were evaluated more positively in the novel context than in the second learning context, but this contrast failed to reach statistical significance, χ^2 (1) = 1.80, p = .180, r = 0.10. Conversely, negative-positive targets were evaluated more negatively in the first learning context as compared to the second learning context, χ^2 (1) = 12.45, $p < .001$, r = 0.26. Moreover, negative-positive targets were evaluated more negatively in the novel context than in the second learning context, but this contrast also failed to reach statistical significance, χ^2 (1) = 1.23, p = .266, r = 0.08.

3.2.2. NFS

The same analyses were repeated for NFS (see Table 2). In the full LME model, the three-way interaction between Valence Order, Measurement Context, and NFS was not significant. In the reduced model, neither the two-way interaction between Valence Order and NFS nor the two –way interaction between measurement Context and NFS was significant. The final model revealed a significant main effect of Valence Order, which was qualified by a significant two-way interaction between Valence Order Measurement Context. The pairwise comparisons decomposing this interaction suggested a pattern of

[5] A sample size of 184 provides 80% power to detect a relation of r = 0.20 between the two components of contextualized attitude change and individual difference measures.

[6] Correlation analyses with the three individual difference measures revealed significant positive correlations between PFC and NFS (r = 0.51, $p < .001$), between PFC and ITP (r = 0.28, $p < .001$), and between NFS and ITP (r = 0.26, $p < .001$).

Table 2
LME model results and Bayes factors, Experiment 2.

Model	Predictor	Model *df*	Residual *df*	*F*	*p*	R_p^2	BF_{10}	BF_{10} interpretation
PFC	Valence Order	1	910	20.41	< .001	0.022	1400.56	Decisive evidence for H_1
	Measurement Context	2	910	0.90	.406	0.002	0.02	Very strong evidence for H_0
	PFC	1	181	0.96	.328	0.006	0.14	Substantial evidence for H_0
	Valence Order × Measurement Context	2	910	7.70	< .001	0.017	33.22	Very strong evidence for H_1
	Valence Order × PFC	1	907	0.26	.612	0.000	0.12	Substantial evidence for H_0
	Measurement Context × PFC	2	907	2.83	.059	0.006	0.22	Substantial evidence for H_0
	Valence Order × Measurement Context × PFC	2	905	0.10	.904	0.000	0.04	Strong evidence for H_0
NFS	Valence Order	1	910	20.41	< .001	0.022	1400.56	Decisive evidence for H_1
	Measurement Context	2	910	0.90	.406	0.002	0.02	Very strong evidence for H_0
	NFS	1	181	0.01	.915	0.000	0.09	Strong evidence for H_0
	Valence Order × Measurement Context	2	910	7.70	< .001	0.017	33.12	Very strong evidence for H_1
	Valence Order × NFS	1	907	2.17	.141	0.002	0.31	Substantial evidence for H_0
	Measurement Context × NFS	2	907	0.30	.738	0.001	0.02	Very strong evidence for H_0
	Valence Order × Measurement Context × NFS	2	905	0.30	.739	0.001	0.04	Strong evidence for H_0
ITP	Valence Order	1	910	20.41	< .001	0.022	1400.56	Decisive evidence for H_1
	Measurement Context	2	910	0.90	.406	0.002	0.02	Very strong evidence for H_0
	ITP	1	181	0.20	.655	0.001	0.10	Substantial evidence for H_0
	Valence Order × Measurement Context	2	910	7.70	< .001	0.017	33.18	Very strong evidence for H_1
	Valence Order × ITP	1	907	2.29	.130	0.003	0.32	Substantial evidence for H_0
	Measurement Context × ITP	1	907	1.37	.255	0.003	0.05	Strong evidence for H_0
	Valence Order × Measurement Context × ITP	2	905	0.46	.628	0.001	0.05	Strong evidence for H_0

Note. LME analyses were conducted using the lmerTest package (Kuznetsova et al., 2016) in RStudio (RStudio Team, 2016) using Kenward-Roger approximated degrees of freedom. Residual degrees of freedom for each effect reflect removal of non-significant higher order effects, which were removed from the model for a more accurate estimation of lower order interactions and main effects (e.g., Kutner et al., 2005). Bayes factors were obtained using the BayesFactor package (Morey & Rouder, 2015) in RStudio (RStudio Team, 2016). To obtain the Bayes factor, a model including only lower level effects (or an empty model for the main effects) was compared to a model containing both the lower level effects and the effect of interest. Evidence category labels for Bayes Factors follow recommendations from Wetzles and Wagenmakers (2012).

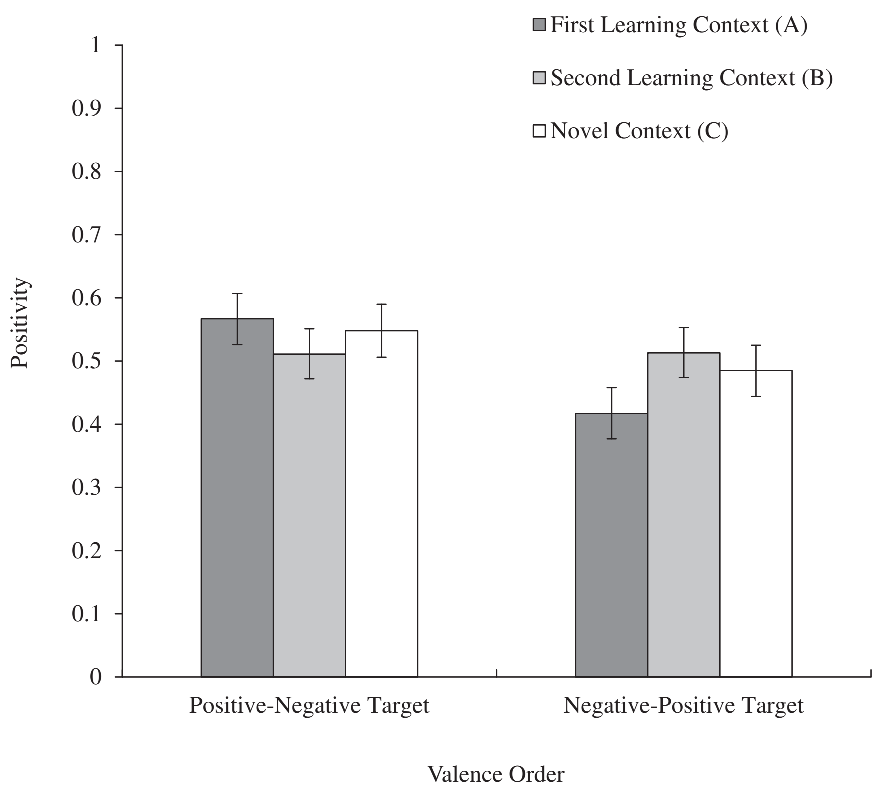

Fig. 3. Target evaluations as a function Valence Order (positive-negative vs. negative-positive) and Measurement Context, Experiment 2. Higher scores reflect more positive evaluations. Error bars represent 95% confidence intervals. Values represent descriptives at the group level collapsed across levels of PFC, NFS, and ITP.

contextualized attitude change in both Valence Order conditions (see Fig. 3). For positive-negative targets, evaluations were more positive in the first learning context as compared to the second learning context, $\chi^2(1) = 4.04$, $p = .044$, $r = 0.15$. Moreover, positive-negative targets were evaluated more positively in the novel context than in the second learning context, but this contrast failed to reach statistical significance, $\chi^2(1) = 1.80$, $p = .180$, $r = 0.10$. Conversely, negative-positive targets were evaluated more negatively in the first learning context as compared to the second learning context, $\chi^2(1) = 12.45$, $p < .001$, $r = 0.26$. Moreover, negative-positive targets were evaluated more negatively in the novel context than in the second learning context, but this contrast also failed to reach statistical significance, χ^2 (1) = 1.23, p = .266, r = 0.08.

3.2.3. ITP

Finally, the LME model was repeated for ITP (see Table 2). As in the other two models, the three-way interaction between Valence Order, Measurement Context, and ITP was not significant. This interaction was removed from the model, and the model was re-estimated. In the reduced model, neither the two-way interaction between Valence Order and ITP nor the two-way interaction between Measurement Context and ITP was significant. These interactions were, thus, removed from the model, and the model was re-estimated. The final model revealed a significant main effect of Valence Order, which was qualified by a significant two-way interaction between Valence Order and Measurement Context. The pairwise comparisons decomposing this interaction suggested a pattern of contextualized attitude change in both Valence Order conditions (see Fig. 3). For positive-negative targets, evaluations were more positive in the first learning context as compared to the second learning context, χ^2 (1) = 4.04, p = .044, r = 0.15. Moreover, positive-negative targets were evaluated more positively in the novel context than in the second learning context, but this contrast failed to reach statistical significance, χ^2 (1) = 1.80, p = .180, r = 0.10. Conversely, evaluations of negative-positive targets were more negative in the first learning context as compared to the second learning context, χ^2 (1) = 12.45, $p < .001$, r = 0.26. Moreover, negative-positive targets were evaluated more negatively in the novel context than in the second learning context, but this contrast also failed to reach statistical significance, χ^2 (1) = 1.23, p = .266, r = 0.08.

3.2.4. Bayesian analysis

As in Experiment 1, we also conducted Bayesian analyses for all effects in each of the three models (see Table 2). In line with the LME analyses reported above, these analyses provided very strong evidence for contextualized attitude change, as captured by the two-way interaction between Valence Order and Measurement Context. Moreover, there was substantial to strong evidence that the two-way interaction between Valence Order and Measurement Context was not qualified by higher-order interactions with PFC, NFS, or ITP.

3.3. Discussion

The goal of Experiment 2 was to replicate the findings of Experiment 1 with a larger sample using the same materials. Although some of the relevant contrasts revealed weaker (and non-significant) effects in the current study, evaluations were more strongly influenced by newly learned counterattitudinal information in the context in which the counterattitudinal information was learned, replicating the findings of Experiment 1 and earlier research using the same paradigm (e.g., Gawronski et al., 2014). Further, contextualized attitude change was not moderated by PFC, NFS, and ITP. Although this result was counter to our hypotheses, it replicates the results of Experiment 1.

4. Experiment 3

Our findings thus far suggest that individual differences in PFC, NFS, and ITP do not moderate contextualized attitude change. Based on the assumptions of Gawronski et al.'s (2010) representational theory, this conclusion further suggests that individual differences along the three dimensions do not moderate the amount of attention paid to the context following an expectancy-violation. Experiment 3 aimed to provide further evidence for our conclusions by directly testing whether individual differences in PFC, NFS, and ITP influence attention to context during the encoding of expectancy-violating information. To the extent that the assumptions of the representational theory are correct, the null effects of the three individual difference variables obtained for contextualized attitude change should generalize to attention to context following expectancy-violation. Yet, if individual differences along the three dimensions influence attention in the manner we initially assumed in the derivation of our predictions, such a finding would pose a challenge to Gawronski et al.'s (2010) assumptions about the mechanisms underlying contextualized attitude change. In the latter case, the theory would have to explain why individual differences in attention to context following an expectancy-violation do not translate into corresponding differences in contextualized attitude change.

To address this question, participants in Experiment 3 completed the same individual difference measures as in Experiments 1 and 2. Participants were then asked to form an impression of a male individual on the basis of 30 statements describing his behavior. Adopting a paradigm by Gawronski et al. (2014), each statement was presented individually against one of ten different colored backgrounds. The first 20 statements described either positive or negative behaviors performed by the individual. The 21st statement—which served as the critical target statement—described a behavior that was either congruent or incongruent with the valence of the initial 20 statements. Finally, participants viewed nine additional statements that were consistent with the valence of the initial 20 statements. After the impression formation task, participants completed a surprise recognition test in which they were asked to indicate the background color against which the critical target statement was presented. Based on evidence that expectancy-violations enhance attention to context (see Roese & Sherman, 2007), memory for the background color of the target statement served as the primary dependent variable (e.g., Brannon et al., 2017; Brannon & Gawronski, in press; Gawronski et al., 2014; Ye et al., 2017). Our main question was whether the memory advantage for the background color of expectancy-incongruent statements (compared to expectancy-congruent statements) varies as a function of PFC, NFS, and ITP.

4.1. Method

4.1.1. Participants and design

One-hundred-and-ninety-one psychology undergraduates completed the study in exchange for course credit. Due to experimenter error, full data are not available for four participants. For two additional participants, the participant number was duplicated, such that data from different components of the study could not be reconciled. This error resulted in the exclusion of two additional participants, leaving us with a final sample of 185 participants (113 women, 72 men; M_{age} = 19.20 years, SD_{age} = 1.11 years).[7] Participants were randomly assigned to the four conditions of a 2 (Impression Valence: positive vs. negative) × 2 (Target Valence: positive vs. negative) between-subjects design.

4.1.2. Individual difference measures

As in Experiments 1 and 2, participants first completed individual difference measures of PFC, NFS, and ITP. The three measures had good reliabilities with Cronbach's αs of 0.87, 0.81, and 0.88, respectively.

4.1.3. Impression formation task

After completing the individual difference measures, participants were asked to form an impression of a target individual by reading statements about behaviors performed by that individual. Each statement was presented beneath the individual's picture for 5000 ms, with each statement-picture pair being presented against one of ten different background colors. Participants were presented with a total of 30 statements in the followings sequence: 20 initial impression statements,

[7] The desired sample size was 200 participants, but we were only able to recruit 191 participants before data collection had to stop at the end of the semester. Based on effect sizes in previous research using the same paradigm (e.g., Brannon et al., 2017), the final sample of 185 provides a power of 96% to detect an expectancy-violation effect, and a power of 80% to detect a three-way interaction between impression valence, target valence, and a given individual difference variable with an odds ratio of 0.65 or 1.54 (depending on the direction of the effect).

Table 3
Logistic regression results and Bayes factors, Experiment 3.

Model	Predictor	*B*	*SE*	Wald *Z*	*p*	*OR*	BF_{10}	BF_{10} interpretation
PFC	Impression Valence	0.05	0.32	0.02	.889	1.05	0.12	Substantial evidence for H_0
	Target Valence	−0.48	0.33	2.19	.139	0.62	0.35	Anecdotal evidence for H_0
	PFC	−0.05	0.16	0.11	.736	0.95	0.06	Strong evidence for H_0
	Impression Valence × Target Valence	−2.39	0.71	11.42	.001	0.09	178.40	Decisive evidence for H_1
	Impression Valence × PFC	0.33	0.36	0.83	.363	1.39	0.55	Anecdotal evidence for H_0
	Target Valence × PFC	−0.03	0.36	0.01	.930	0.97	0.32	Substantial evidence for H_0
	Impression Valence × Target Valence × PFC	0.68	0.74	0.84	.361	1.97	0.36	Anecdotal evidence for H_0
NFS	Impression Valence	0.04	0.32	0.02	.890	1.05	0.12	Substantial evidence for H_0
	Target Valence	−0.48	0.33	2.15	.143	0.62	0.36	Anecdotal evidence for H_0
	NFS	−0.05	0.24	0.05	.827	0.95	0.06	Strong evidence for H_0
	Impression Valence × Target Valence	−2.46	0.71	12.04	.001	0.09	180.94	Decisive evidence for H_1
	Impression Valence × NFS	−0.02	0.55	0.00	.972	0.98	0.30	Substantial evidence for H_0
	Target Valence × NFS	0.00	0.55	0.00	.996	1.00	0.28	Substantial evidence for H_0
	Impression Valence × Target Valence × NFS	−0.24	1.10	0.05	.827	0.79	0.50	Anecdotal evidence for H_0
ITP	Impression Valence	0.09	0.33	0.07	.795	1.09	0.12	Substantial evidence for H_0
	Target Valence	−0.47	0.32	2.12	.146	0.62	0.35	Anecdotal evidence for H_0
	ITP	0.14	0.19	0.55	.459	1.15	0.08	Strong evidence for H_0
	Impression Valence × Target Valence	−2.46	0.71	12.11	.001	0.09	168.95	Decisive evidence for H_1
	Impression Valence × ITP	0.14	0.40	0.12	.730	1.15	0.22	Substantial evidence for H_0
	Target Valence × ITP	−0.04	0.40	0.01	.919	0.96	0.23	Substantial evidence for H_0
	Impression Valence × Target Valence × ITP	0.06	0.80	0.00	.944	1.06	0.36	Anecdotal evidence for H_0

Note. Logistic regression NHSTs were conducted using SPSS. For proper estimation of each effect, main effects, two-way interactions, and three-way interactions were entered into separate blocks. Bayes factors were obtained by first running a Bayesian logistic regression via the stan_glm function in the rstanarm package (Stan Development Team, 2016) in RStudio (RStudio Team, 2016). To obtain a Bayes factor for each effect, model comparisons were then conducted using the bridgesampling package (Gronau & Singmann, 2017) in RStudio.

1 target statement, and 9 filler statements. For half of the participants, the initial 20 statements and the 9 filler statements depicted positive behaviors; for the other half the initial 20 statements and the 9 filler statements depicted negative behaviors. Orthogonal to the manipulation of impression valence, the valence of the target statement was positive for half of the participants and negative for the other half. Hence, the valence of target statement could be either congruent or incongruent with the valence of the initial impression statements. The background colors were randomized in a blocked manner, such that each color appeared once during the first block of 10 statements, once during the second block of 10 statements, and once during the third block of 10 statements. The critical target statement was presented against the same background color in each of the four experimental conditions (i.e., blue).

4.1.4. Surprise recognition test

After the impression formation task, participants completed a surprise recognition test in which they were asked to indicate the background color against which a given statement was presented during the impression formation task. The recognition test included a total of seven statements in the following order: three statements randomly selected from the impression and filler statements, the critical target statement, and three additional statements randomly selected from the impression and filler statements. Each statement was presented against a black background below 10 colored squares (one for each of the 10 background colors used in the impression formation task) numbered from 0 to 9. Participants were asked to indicate the background color of the presented statement in the impression formation task by pressing the key on the keyboard corresponding to the number. Based on previous research showing a memory advantage for the background color for expectancy-disconfirming information compared to expectancy-confirming information (e.g., Brannon et al., 2017; Brannon & Gawronski, in press; Gawronski et al., 2014; Ye et al., 2017), this memory advantage served as an indicator for enhanced attention to context during the encoding of expectancy-violating information.

4.2. Results

To test whether any of the three individual difference variables moderated attention to context in response to inconsistency, three separate logistic regressions were conducted with Impression Valence and Target Valence as dummy-coded predictors, one of the three individual difference measures as continuous predictor, and memory for the background color of the target statement as the outcome.[8]

A model containing Impression Valence, Target Valence, and PFC as predictors revealed a significant interaction between Impression Valence and Target Valence (see Table 3). Replicating past research (e.g., Brannon et al., 2017; Brannon & Gawronski, in press; Gawronski et al., 2014; Ye et al., 2017), participants better remembered the background color of the target statement when a positive target statement followed negative impression statements than when a positive target statement followed positive impression statements, $B = 1.34$, $SE = 0.54$, Wald $Z = 6.28$, $p = .012$, $OR = 3.82$. Conversely, participants better remembered the background color of the target statement when a negative target statement followed positive impression statements than when a negative target statement followed negative impression statements, $B = 1.21$, $SE = 0.45$, Wald $Z = 6.12$, $p = .013$, $OR = 3.07$ (see Fig. 4). This interaction, however, was not qualified by a higher-order interaction with PFC.

Similarly, a model containing Impression Valence, Target Valence, and NFS as predictors revealed a significant two-way interaction between Impression Valence and Target Valence (see Table 3). As for the model containing PFC, participants better remembered the background color of the target statement when a positive target statement followed negative impression statements than when a positive target statement followed positive impression statements, $B = 1.34$, $SE = 0.54$, Wald $Z = 6.28$, $p = .012$, $OR = 3.82$. Conversely, participants better remembered the background color of the target statement when a negative target statement followed positive impression statements than when a negative target statement followed negative impression statements, $B = 1.22$, $SE = 0.45$, Wald $Z = 6.12$, $p = .013$, $OR = 3.07$ (see Fig. 4). This interaction, however, was not qualified by a higher-order interaction with NFS.

[8] Correlation analyses with the three individual difference measures revealed significant positive correlations between PFC and NFS ($r = 0.55$, $p < .001$), between PFC and ITP ($r = 0.15$, $p = .042$) and between NFS and ITP ($r = 0.14$, $p = .050$).

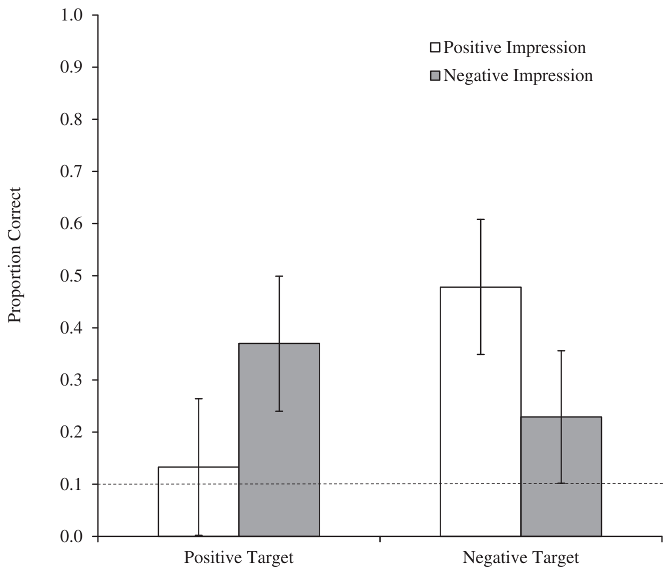

Fig. 4. Mean proportion of correct background recognition as a function of target statement valence and initial impression valence, Experiment 3. Dotted line represents chance responding. Error bars represent 95% confidence intervals. Values represent descriptives at the group level collapsed across levels of PFC, NFS, and ITP.

Finally, when ITP was included as a continuous predictor, there was again a significant two-way interaction between Impression Valence and Target Valence (see Table 3). As in the previous models, participants better remembered the background color of the target statement when a positive target statement followed negative impression statements than when a positive target statement followed positive impression statements, $B = 1.30$, $SE = 0.54$, Wald $Z = 5.83$, $p = .016$, $OR = 3.67$. Conversely, participants better remembered the background color of the target statement when a negative target statement followed positive impression statements than when a negative target statement followed negative impression statements, $B = 1.15$, $SE = 0.46$, Wald $Z = 6.36$, $p = .012$, $OR = 3.15$ (see Fig. 4). This interaction, however, was not qualified by a higher-order interaction with ITP.

As in Experiments 1 and 2, we also conducted Bayesian analyses for the effects in each of the three models (see Table 3). These analyses provided decisive evidence for the memory advantage for the background color of expectancy-incongruent statements (compared to expectancy-congruent statements), captured by the two-way interaction between Impression Valence and Target Valence. Moreover, there was anecdotal evidence that this interaction was not qualified by higher-order interactions with PFC, NFS, or ITP. Although the evidence against moderation by the three individual difference variables is weaker in this experiment, it is still in line with the results of Experiments 1 and 2.

4.3. Discussion

Experiment 3 investigated whether individual differences in PFC, NFS, and ITP moderate attention to context following an expectancy-violation. To the extent that the assumptions of Gawronski et al.'s (2010) representational theory are correct, the null effects for the hypothesized relation between the three individual difference variables and contextualized attitude change should generalize to attention to context following expectancy-violations. Yet, if individual differences along the three dimensions influence attention in the manner we initially proposed in the derivation of our predictions, such a finding would pose a challenge to Gawronski et al.'s (2010) assumptions about the mechanisms underlying contextualized attitude change.

Replicating earlier research using the same paradigm (Brannon et al., 2017; Brannon & Gawronski, in press; Gawronski et al., 2014; Ye et al., 2017), participants showed a memory advantage for the background color of expectancy-disconfirming information compared to expectancy-confirming information. However, this memory advantage remained unqualified by individual differences in PFC, NFS, and ITP. This finding disconfirms our initial hypothesis that individual differences along the three dimensions would moderate attention to context in response to expectancy-violating information. Yet, it reconciles our conclusion from Experiments 1 and 2 with Gawronski et al.'s (2010) representation theory. Together, the results of the three studies suggest that individual differences in PFC, NFS, and ITP do not moderate attention to context following an expectancy-violation and, thus, the emergence of contextualized attitude change.

5. General discussion

The goal of the current research was to investigate whether contextualized attitude change is moderated by individual differences in PFC, NFS, and ITP. We hypothesized that the three individual difference dimensions would moderate contextualized attitude change by influencing the amount of attention paid to the context following an expectancy-violation. Contrary to our hypotheses, none of the three individual difference variables moderated patterns of contextualized attitude change found in past research (Experiments 1 and 2). This conclusion was supported by the results of both NHSTs and Bayesian analyses.[9] There was also no evidence that the three individual difference variables moderated attention to context during exposure to expectancy-violating information (Experiment 3). Taken together, these results disconfirm our hypotheses regarding the roles of PFC, NFS, and ITP in contextualized attitude change. Instead, they suggest that contextualized attitude change is a general phenomenon that is relatively robust across individuals.

5.1. Implications for contextualized attitude change

In a recent meta-analysis, Gawronski et al. (2015) noted that research under the framework of Gawronski et al.'s (2010) representational theory revealed mixed results. Although meta-analytic effect sizes of the two critical comparisons (see Fig. 1) were significantly larger than zero, effect sizes were relatively small (*d*s = 0.25 and 0.17) and varied considerably across studies. In light of these findings, Gawronski et al. (2015) sought to identify potential boundary conditions that may account for the observed heterogeneity in effect sizes. In their meta-analysis, effect sizes depended on (1) attention to context, (2) context-valence contingencies during the learning of evaluative information, and (3) the country in which a given study was conducted. An additional boundary condition has been identified by Brannon and Gawronski (2017), who found that (1) counterattitudinal information that is highly diagnostic as well as (2) new information that changes the meaning of earlier information can lead to changes in evaluations that are independent of the context.

The current research expands on these findings by investigating the potential role of individual differences in PFC, NFS, and ITP. To the extent that some people do not show enhanced attention to context following an expectancy-violation, these individuals should be less susceptible to show contextualized attitude change according to Gawronski et al.'s (2010) representational theory. Hence, the small effect sizes in Gawronski et al.'s (2015) meta-analysis may at least in part be due to variations across individuals, in that some people may show strong effects of contextualized attitude change, whereas others may

[9] To obtain greater statistical power for detecting small effects, we also conducted a combined analysis of the data from Experiments 1 and 2 (see Supplementary materials). The NHSTs and Bayesian analyses of the combined data further support our conclusions that contextualized attitude change is not moderated by individual differences in PFC, NFS or ITP.

show no evidence for contextualized attitude change at all. Although we cannot rule out that other trait dimensions contribute to variations in contextualized attitude change across individuals, the current findings suggest that contextualized attitude change is unaffected by individual differences in PFC, NFS, and ITP.

Indeed, counter to our original goal to identify boundary conditions, the current findings provide further evidence for the generality of contextualized attitude change. In this sense, our findings expand on other research that demonstrated the generality of contextualized attitude change despite the initial aim to identify boundary conditions. For example, based on evidence that individuals in Eastern cultures tend to have a higher tolerance for inconsistency than individuals in Western cultures (Peng & Nisbett, 1999; Spencer-Rodgers, Williams, & Peng, 2010), Ye et al. (2017) investigated whether contextualized attitude change replicates in participants from Eastern cultures. Counter to the prediction that contextualized attitude change is limited to Western cultures, Ye et al. replicated the typical pattern of contextualized attitude change in participants from Singapore. In a follow-up study, the authors further showed that both Canadian and Singaporean participants showed enhanced attention to context following an expectancy-violation, similar to the findings of the current studies. Together with research showing similar patterns in animal learning (see Bouton, 2004), these findings collectively support the generality of contextualized attitude change across individuals, cultures, and species.

5.2. Implications for PFC, NFS, and ITP

In addition to demonstrating the generality of contextualized attitude change, the current research also has important implications for individual differences in PFC, NFS, and ITP. Specifically, the results of the current research support a more nuanced understanding of these individual difference variables and their impact on the processing of social information. Although our findings may seem surprising in light of past research showing that individual differences in PFC, NFS, and ITP influence people's responses to inconsistency (e.g., Cialdini et al., 1995; Gocłowska et al., 2014; Newby-Clark et al., 2002; Plaks et al., 2001), an important difference between the current work and earlier research is the stage of inconsistency processing that is investigated. Because of this critical difference, the current research is not at odds with past research on PFC, NFS, and ITP but instead highlights the need for a more nuanced view of the three individual difference constructs and their psychological effects.

In line with this argument, Gawronski and Brannon (in press) suggest that research on cognitive consistency often conflates three stages of processing: (1) the identification of an inconsistency, (2) the feelings experienced in response to inconsistency, and (3) the reconciliation of inconsistency (see also Gawronski, Peters, & Strack, 2008). Past research on PFC, NFS, and ITP has primarily focused on the latter two stages. For example, Newby-Clark et al. (2002) found that individuals high in PFC experienced greater arousal in response to ambivalence than individuals low in PFC, reflecting a difference at the second stage of inconsistency processing. Reflecting a difference at the third stage of inconsistency processing, PFC has also been found to moderate attitude change in response to inconsistencies between attitudes and behaviors (Cialdini et al., 1995), and NFS has been shown to moderate cognitive flexibility and performance following an inconsistency (Gocłowska et al., 2014). Further, individuals high in NFS are more likely to integrate information that is consistent with their social expectancies into an impression than those low in NFS (Thompson, Roman, Moskowitz, Chaiken, & Bargh, 1994), again reflecting a difference at the third stage.

In contrast to these findings, the current research focused on the first stage of inconsistency processing: the identification of inconsistency. According to Gawronski et al.'s (2010) representational theory, contextualized attitude change is a product of attentional processes during the encoding of expectancy-violating information. These attentional processes are assumed to arise as a result of an inconsistency being identified, which signals an error in one's system of beliefs (Gawronski, 2012; Gawronski & Brannon, in press). From this perspective, the current results suggest that individual differences in PFC and NFS do not moderate the identification of inconsistencies and the attentional processes resulting from identified inconsistency. Nevertheless, PFC and NFS may moderate the level of arousal elicited by inconsistency (Newby-Clark et al., 2002) as well as the reconciliation of inconsistency (Cialdini et al., 1995). Thus, in addition to supporting the generality of contextualized attitude change, our findings help to better understand the impact of PFC and NFS on social information processing by specifying the processing stages that are influenced by these personality traits.

Although the distinction between different processing stages reconciles the current findings with previous research on PFC and NFS, a similar conclusion may still seem at odds with previous research on ITP (Plaks et al., 2001, 2005). For example, counter to the conclusion that ITP do not influence attentional processes at the first processing stage, Plaks et al. (2001) suggest that entity theorists pay greater attention to expectancy-*consistent* information whereas incremental theorists pay greater attention to expectancy-*inconsistent* information. Closer examination of the paradigms used in this research, however, suggests that our results do not necessarily conflict with their results. First, the majority of studies by Plaks et al. (2001) examined what information people attend to when there are competing demands. For example, when *both* expectancy-consistent and expectancy-inconsistent information is simultaneously available, entity theorists pay more attention to the former while incremental theorists pay more attention to the latter. In contrast to the forced choice paradigm in Plaks et al.'s research, participants in our studies were presented with either expectancy-consistent or expectancy-inconsistent information. Second, Plaks et al.'s results do not speak to whether entity and incremental theorists identify inconsistencies differently. Instead, they suggest that, once an inconsistency has been identified, entity versus incremental theorists may process different pieces of information more deeply in order to maintain their beliefs about the stability versus malleability of personality (see also Plaks et al., 2005). Thus, their findings speak to the processes involved in the reconciliation of an inconsistency (Stage 3) rather than the identification of inconsistency (Stage 1). From this perspective, our findings help to better understand the impact of ITP on social information processing by providing deeper insights into the processing stages that are influenced by ITP.

Although Gawronski and Brannon's (in press) three-stage model reconciles the current results with past work, this model has not yet been formally tested. Thus, an interesting direction for future research would be to examine the effects of PFC, NFS, and ITP on the identification of inconsistency, affective responses to inconsistency, and the resolution of inconsistency. Research on these questions may provide both (1) valuable evidence in support of the three-stage model of inconsistency and (2) more nuanced insights into the psychological effects of PFC, NFS, and ITP.

The results of the current research additionally contribute to a growing body of research on the temporal dynamics of surprise (Noordewier & Breugelmans, 2013; Noordewier, Topolinski, & Van Dijk, 2016). To reconcile discrepant findings regarding the subjective experience of surprise, Noordewier et al. (2016) suggest that surprising events tend to be experienced as aversive immediately following their identification (see also Topolinski & Strack, 2015). However, once perceivers had an opportunity to make sense of the unexpected event, their affective responses may reflect the hedonic valence of the observed event, such that positive events elicit positive affect and negative events elicit negative affect (see also Noordewier & Breugelmans, 2013). For example, a student might initially experience negative arousal when they unexpectedly receive an “A” on an exam. Yet, once this individual had sufficient time to make sense of the unexpected outcome, the initial negative affect will be replaced by the positive affect elicited by the positive quality of the unanticipated outcome.

Building on the temporal dynamics of surprise, the current research further suggests that the initial reactions caused by unexpected events may be rather general and unrelated to individual differences in PFC, NFS, and ITP. In line with this conclusion, Gocłowska, Baas, Elliot, and De Dreu (2017) found that participants found schema-violations surprising, regardless of individual differences in NFS and openness to new experiences. Thus, these personality traits seem to influence psychological outcomes via sense-making processes by influencing either (1) how an unexpected event is appraised as time passes or (2) the length of time it takes for someone to make sense of the inconsistency (or both). Although this interpretation would be consistent with our results, our data do not directly speak to this conclusion. Thus, investigating the impact of PFC, NFS, and ITP on the temporal dynamics of inconsistency processing would be a fruitful direction for future research.

5.3. Potential objection

One potential objection against our conclusions is that individual differences in PFC, NFS, and ITP did not moderate contextualized attitude change because participants were not sufficiently invested in the task. In line with this concern, classic theory and research suggests that dissonance effects depend on the subjective importance or self-relevance of the involved cognitions (e.g., Aronson, 1968, 1969; Festinger, 1957; Steele & Liu, 1983). Thus, to the extent that participants in the current studies did not care about the information in the impression formation task, there may have been no arousal in response to inconsistent information and, thus, no attempt to reconcile the inconsistency. However, there are at least two pieces of evidence that argue against such an interpretation of the current findings. First, the view that high subjective importance or self-relevance are necessary to produce responses to inconsistency has been disputed. For example, Proulx et al. (2012) argued that inconsistencies per se are sufficient to produce responses aimed toward reconciliation, regardless of their importance to an individual. In line with this assumption, some studies have found responses to inconsistencies as mundane as anomalous playing cards (e.g., Sleegers, Proulx, & van Beest, 2015). Second, although the three individual difference variables did not show any effects in the current studies, we did find significant evidence for contextualized attitude change and enhanced attention to the background against which inconsistent information was presented (see also Brannon et al., 2017; Brannon & Gawronski, in press; Gawronski et al., 2014; Ye et al., 2017). If high subjective importance or self-relevance are necessary to produce responses to inconsistency and participants did not care about the presented information in the current studies, there should be no evidence for enhanced attention to the context of expectancy-violating information and no evidence for contextualized attitude change. Consequently, the current findings cannot be explained by lack of involvement or investment in the task.

6. Conclusion

The current research sought to investigate whether contextualized attitude change depends on individual differences in PFC, NFS, and ITP. Contrary to our hypotheses, none of the three individual difference variables moderated (1) contextualized attitude change and (2) attention to context following an expectancy-violation. Nevertheless, the current research replicated past findings on contextualized attitude change and attention to context following an expectancy-violation. As such, our findings provide further evidence for the generality of contextualized attitude change across individuals, adding to the growing body of research showing similar effects in animals (Bouton, 2004) and across different cultures (Ye et al., 2017). Additionally, the current research provides valuable insights for research on cognitive consistency by specifying the processing stages that are influenced by individual differences in PFC, NFS, and ITP.

Open practices

The materials, data, and analysis files for the studies reported in this article are publicly available at https://osf.io/usgz3/. The level of detail provided in the analysis files is sufficient for independent researchers to reproduce the results reported in the article using the data files available on the repository. Additionally, sufficient detail is provided in the materials files for independent researchers to replicate the reported studies or to implement the procedures in new research.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jesp.2018.03.015.

References

Aronson, E. (1968). Dissonance theory: Progress and problems. In R. P. Abelson, E. Aronson, W. J. McGuire, T. M. Newcombe, M. J. Rosenberg, & P. H. Tannenbaum (Eds.). *Theories of cognitive consistency: A sourcebook* (pp. 5–27). Skokie, IL: Rand-McNally.

Aronson, E. (1969). The theory of cognitive dissonance: A current perspective. *Advances in Experimental Social Psychology, 4*, 1–34.

Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning and Memory, 11*, 485–494.

Brannon, S. M., & Gawronski, B. (2017). A second chance for first impressions? Exploring the context (in)dependent updating of implicit evaluations. *Social Psychological and Personality Science, 8*, 275–283.

Brannon, S. M., & Gawronski, B. (in press). In search of a negativity bias in expectancy-violation. Social Cognition.

Brannon, S. M., Sacchi, D. L. M., & Gawronski, B. (2017). (In)consistency in the eye of the beholder: The roles of warmth, competence, and valence in lay perceptions of inconsistency. *Journal of Experimental Social Psychology, 70*, 80–94.

Chiu, C.-Y., Hong, Y.-Y., & Dweck, C. S. (1997). Lay dispositionism and implicit theories of personality. *Journal of Personality and Social Psychology, 73*, 19–30.

Cialdini, R. B., Trost, M. R., & Newsom, J. T. (1995). Preference for consistency: The development of a valid measure and the discovery of surprising behavioral implications. *Journal of Personality and Social Psychology, 69*, 318–328.

De Rosario-Martinez, H. (2015). phia: Post-Hoc interaction analysis. R package version 0.2-1. https://CRAN.R-project.org/package=phia.

Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological Review, 95*, 256–273.

Festinger, L. (1957). *A theory of cognitive dissonance.* Evanston, IL: Row Peterson.

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

Gawronski, B. (2012). Back to the future of dissonance theory: Cognitive consistency as a core motive. *Social Cognition, 30*, 652–668.

Gawronski, B., & Brannon, S. M. (in press). What is cognitive consistency and why does it matter? In E. Harmon-Jones (Ed.), Cognitive dissonance: Progress on a pivotal theory in social psychology (2nd edition). Washington, DC: American Psychological Association.

Gawronski, B., & Cesario, J. (2013). Of mice and men: What animal research can tell us about context effects on automatic responses in humans. *Personality and Social Psychology Review, 17*, 187–215.

Gawronski, B., Hu, X., Rydell, R. J., Vervliet, B., & De Houwer, J. (2015). Generalization and contextualization in automatic evaluation revisited: A meta-analysis of successful and failed replications. *Journal of Experimental Psychology: General, 144*, e50-e64.

Gawronski, B., Peters, K. R., & Strack, F. (2008). Cross-cultural differences vs. universality in cognitive dissonance: A conceptual reanalysis. In R. M. Sorrentino, & S. Yamaguchi (Eds.). *Handbook of motivation and cognition across cultures* (pp. 297–314). New York: Elsevier.

Gawronski, B., Rydell, R. J., De Houwer, J., Brannon, S. M., Ye, Y., Vervliet, B., & Hu, X. (2018). Contextualized attitude change. *Advances in Experimental Social Psychology, 57*, 1–52.

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General, 139*, 683–701.

Gawronski, B., Ye, Y., Rydell, R. J., & De Houwer, J. (2014). Formation, representation, and activation of contextualized attitudes. *Journal of Experimental Social Psychology, 54*, 188–203.

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

Gocłowska, M. A., Baas, M., Crisp, R. J., & De Dreu, C. K. W. (2014). Whether social schema violations help or hurt creativity depends on need for structure. *Personality and Social Psychology Bulletin, 40*, 959–971.

Gocłowska, M. A., Baas, M., Elliot, A. J., & De Dreu, C. K. W. (2017). Why schema-violations are sometimes preferable to schema-consistencies: The role of interest and openness to experience. *Journal of Research in Personality, 66*, 54–69.

Gronau, Q. F., & Singmann, H. (2017). Bridgesampling: Bridge sampling for marginal likelihoods and Bayes factors. R package version 0.4-0. https://CRAN.R-project.org/

package = bridgesampling.

Kruglanski, A. W. (1990). Lay epistemic theory in social-cognitive psychology. *Psychological Inquiry, 1*, 181–197.

Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: "Seizing" and "freezing". *Psychological Review, 103*, 263–283.

Kutner, M. H., Nachtsheim, C. J., Neter, J., & Li, W. (2005). *Applied linear statistical models* (5th ed.). New York, NY: McGraw-Hill.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). lmerTest: Tests in linear mixed effects models. R package version 2.0-33. https://CRAN.R-project.org/package = lmerTest.

Lipsey, M. W., & Wilson, D. S. (2001). *Practical meta-analysis.* Thousand Oaks, CA: Sage.

Morey, R. D., & Rouder, J. N. (2015). BayesFactor: Computation of Bayes factors for common designs. R package version 0.9.12-2. https://CRAN.R-project.org/package = BayesFactor.

Neuberg, S. L., & Newsom, J. T. (1993). Personal need for structure: Individual differences in the desire for simpler structure. *Journal of Personality and Social Psychology, 65*, 113–131.

Newby-Clark, I. R., McGregor, I., & Zanna, M. P. (2002). Thinking and caring about cognitive inconsistency: When and for whom does attitudinal ambivalence feel uncomfortable? *Journal of Personality and Social Psychology, 82*, 157–166.

Noordewier, M. K., & Breugelmans, S. M. (2013). On the valence of surprise. *Cognition and Emotion, 27*, 1326–1334.

Noordewier, M. K., Topolinski, S., & Van Dijk, E. (2016). The temporal dynamics of surprise. *Social and Personality Compass, 10*, 136–149.

Payne, B. K., Cheng, S. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89*, 277–293.

Peng, K., & Nisbett, R. E. (1999). Culture, dialectics, and reasoning about contradiction. *American Psychologist, 54*, 741–754.

Plaks, J. E. (2017). Implicit theories: Assumptions that shape social and moral cognition. *Advances in Experimental Social Psychology, 56*, 259–310.

Plaks, J. E., Grant, H., & Dweck, C. S. (2005). Violations of implicit theories and the sense of prediction and control: Implications for motivated person perception. *Journal of Personality and Social Psychology, 88*, 245–262.

Plaks, J. E., Stroessner, S. J., Dweck, C. S., & Sherman, J. W. (2001). Person theories and attention allocation: Preferences for stereotypic versus counterstereotypic information. *Journal of Personality and Social Psychology, 80*, 876–893.

Proulx, T., Inzlicht, M., & Harmon-Jones, E. (2012). Understanding all inconsistency-compensation as a palliative response to violated expectations. *Trends in Cognitive Sciences, 16*, 285–291.

Ranganath, K. A., Smith, C. T., & Nosek, B. A. (2008). Distinguishing automatic and controlled components of attitudes form direct and indirect measurement methods. *Journal of Experimental Social Psychology, 44*, 386–396.

Roese, N. J., & Sherman, J. W. (2007). Expectancies. In E. T. Higgins, & A. W. Kruglanski (Eds.). *Social psychology: Handbook of basic principles* (pp. 91–115). (2nd Ed.). New York: Guilford Press.

RStudio Team (2016). RStudio: Integrated development for R. RStudio. *Inc.*Boston: MA URLhttp://www.rstudio.com/.

Rydell, R. J., & Gawronski, B. (2009). I like you, I like you not: Understanding the formation of context-dependent automatic attitudes. *Cognition and Emotion, 23*, 1118–1152.

Sleegers, W. W., Proulx, T., & van Beest, I. (2015). Extremism reduces conflict arousal and increases values affirmation in response to meaning violations. *Biological Psychology, 108*, 126–131.

Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology, 70*, 893–912.

Spencer-Rodgers, J., Williams, M. J., & Peng, K. (2010). Cultural differences in expectations of change and tolerance for contradiction: A decade of empirical research. *Personality and Social Psychology Review, 14*, 296–312.

Stan Development Team (2016). rstanarm: Bayesian applied regression modeling via Stan. R package version 2.13.1. http://mc-stan.org/.

Steele, C. M., & Liu, T. J. (1983). Dissonance processes as self-affirmation. *Journal of Personality and Social Psychology, 45*, 5–19.

Thompson, E. P., Roman, R. J., Moskowitz, G. B., Chaiken, S., & Bargh, J. A. (1994). Accuracy motivation attenuates covert priming: The systematic reprocessing of social information. *Journal of Personality and Social Psychology, 66*, 474–489.

Topolinski, S., & Strack, F. (2015). Corrugator activity confirms immediate negative affect in surprise. *Frontiers in Psychology, 6*, 134.

Wetzles, R., & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin and Review, 19*, 1057–1064.

Ye, Y., Tong, Y.-Y., Chiu, C.-Y., & Gawronski, B. (2017). Attention to context during evaluative learning and context-dependent automatic evaluation: A cross-cultural analysis. *Journal of Experimental Social Psychology, 70*, 1–7.