

IN SEARCH OF A NEGATIVITY BIAS IN EXPECTANCY VIOLATION

Skylar M. Brannon and Bertram Gawronski
University of Texas at Austin

Negative information tends to dominate positive information across many domains. However, previous research did not find any evidence for valence asymmetries in the violation of expectancies. The current research tested whether negativity bias in expectancy violation depends on the amount of prior information that is available about a target. Drawing on extant theories of the negativity bias in impression formation, new negative and new positive information were hypothesized to result in equal expectancy violations when the new information conflicts with large amounts of target-specific information. In contrast, new negative information was hypothesized to result in greater expectancy violations than new positive information when the new information conflicts with small amounts of target-specific information. Three experiments ($N = 972$) disconfirmed the latter hypothesis, showing that new negative and new positive information resulted in equal expectancy violations regardless of the amount of prior information. Implications for the negativity bias, impression formation, and cognitive consistency are discussed.

Keywords: cognitive consistency, expectancy violation, impression formation, negativity bias

The tendency for negative information to dominate positive information has been recognized as one of the most robust phenomena in psychological science (for reviews, see Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Peeters & Czapinski, 1990; Rozin & Royzman, 2001). For example, negative information is given more attention than positive information (e.g., Pratto & John, 1991), negative information tends to influence impressions to a greater extent than positive information (e.g., Kanouse & Hanson, 1972), and negative evaluations form more quickly and

Correspondence concerning this article should be addressed to Skylar M. Brannon, Department of Psychology, University of Texas at Austin, 108 E. Dean Keeton A8000, Austin, TX 78712; E-mail: skylar_brannon@utexas.edu.

are generalized more readily than positive evaluations (e.g., Fazio, Eiser, & Shook, 2004; Shook, Fazio, & Eiser, 2007). Interestingly, these asymmetries occur even when negative and positive information are equal in terms of their overall extremity (see Reeder & Brewer, 1979; Skowronski & Carlston, 1989).

A notable exception to the relatively greater impact of negative over positive information seems to be the violation of expectancies. Past research consistently failed to obtain evidence that expectancy-incongruent negative information elicits a stronger surprise response than expectancy-incongruent positive information. For example, in a study by Cacioppo, Crites, Berntson, and Coles (1993), the strength of neurological markers of expectancy violation did not differ as a function of valence, only as a function of inconsistency. Similarly, in a study by Brannon, Sacchi, and Gawronski (2017), new positive and new negative information was perceived as equally inconsistent with prior impressions of the opposite valence, resulting in equal expectancy violations for the two kinds of information. While these findings may seem surprising in light of evidence for the pervasiveness of the negativity bias, some theories suggest that the emergence of a negativity bias in expectancy violation may depend on the amount of prior information that is available about a specific target. Although these theories do not directly address expectancy violations per se, their postulates suggest that a negativity bias in expectancy violation may occur only when a small amount of target-specific information is available. In contrast, expectancy violations may be equal for new negative and new positive information to the extent that this information conflicts with a large amount of target-specific information.

The current research aimed to test this prediction. Toward this end, we manipulated the amount of prior information participants received about a target person to test whether a small amount of prior information would produce greater expectancy violations for new negative, as compared to new positive, information about that person. Conversely, when participants received a large amount of prior information, we expected to replicate the results of previous studies, showing equal expectancy violations irrespective of the valence of the expectancy-violating information (e.g., Brannon et al., 2017; Cacioppo et al., 1993).

NEGATIVITY BIAS AND EXPECTANCY VIOLATION

Conceptually, two pieces of information are inconsistent to the extent that one piece of information implies the opposite of the other (Festinger, 1957). For example, in the domain of impression formation, two pieces of information about another individual would be inconsistent if one suggests an evaluation that is opposite to the other (e.g., *Bob helped an elderly woman with her groceries* and *Bob kicked a puppy*). Past research suggests that, in cases of such inconsistencies, negative information tends to have a greater impact on overall impressions than positive information. Several theories have been proposed to explain why negative information tends to dominate positive information in the integration of conflicting information. These theories fall into two broad categories: (1) theories maintaining

that negativity biases are due to inherent differences between positive and negative information and (2) theories suggesting that asymmetries between positive and negative information depend on specific informational conditions. Although these theories do not explicitly address the role of valence asymmetries in expectancy violation, their propositions can be used to derive predictions regarding the relative strength of expectancy violations.

Some theories maintain that negativity biases are due to inherent differences between positive and negative information. For example, negative information has been claimed to be inherently more informative for trait inferences than positive information, as negative behaviors are associated with fewer traits than positive behaviors (Birnbaum, 1972; Wyer, 1973). Further, distinct pieces of positive information tend to be more similar to one another than distinct pieces of negative information, and these differences in distinctiveness have been claimed to drive valence asymmetries across many domains (e.g., Alves, Koch, & Unkelbach, 2017). Applied to the current question, the assumption that positive and negative information possess inherently different properties implies that new negative information should produce greater expectancy violations than new positive information under all circumstances. This prediction, however, is difficult to reconcile with previous evidence showing that new negative information and new positive information result in equal expectancy violations (e.g., Brannon et al., 2017; Cacioppo et al., 1993).

Other theories, however, propose that the negativity bias is a conditional phenomenon, which makes it easier to reconcile these theories with past research that failed to find a negativity bias in the violation of expectancies (e.g., Brannon et al., 2017; Cacioppo et al., 1993). Specifically, these theories suggest that the impact of valenced information is not determined in an absolute sense but instead depends on people's broader sets of beliefs and expectancies (for a review, see Skowronski & Carlston, 1989). These theories do not prohibit positive information from being more influential than negative information. Instead, they specify the conditions under which negative information should be more influential than positive information, and vice versa.

For example, expectancy-contrast theories suggest that information receives greater weight depending on its inconsistency with a reference of comparison (e.g., Helson, 1964; Sherif & Sherif, 1967). To the extent that new negative information is contrasted with a large amount of positive information, the negative information will stand out perceptually, receiving greater attention and weighting. Conversely, to the extent that new positive information is contrasted with a large amount of negative information, the positive information will stand out perceptually, receiving greater attention and weighting. Frequency-weight theories similarly suggest that information receives greater weight depending on its statistical infrequency, because infrequent actions are more indicative to underlying traits (Fiske, 1980; Jones & Davis, 1965). From the perspective of these theories, negative information about a person should have a greater impact in contexts where negative information is less frequent than positive information. Conversely, positive information

about a person should have a greater impact in contexts where positive information is less frequent than negative information.

Expectancy-contrast and frequency-weight theories suggest that new negative information and new positive information have the potential to be equally expectancy violating depending on their inconsistency with a reference of comparison. The theories further suggest that this reference of comparison depends on (1) the available information about a focal target and (2) perceivers' default beliefs about other people in general. These two determinants influence perceivers' reference of comparison in a compensatory manner, such that the impact of default beliefs should decrease with increasing amounts of target-specific information. Thus, in conjunction with evidence that default beliefs about other people tend to be positive (i.e., positivity offset; see Cacioppo, Gardner, & Berntson, 1997), expectancy-contrast and frequency-weight theories imply specific predictions about the conditions under which new negative information should lead to stronger expectancy violation than new positive information.

To the extent that a large amount of prior information about a particular target is available, perceivers' reference of comparison should be determined by the available target-specific information with little impact of positive default beliefs. As a result, expectancy violation in response to new information about a target person depends primarily on the (in)consistency of this information with prior target-specific information. For example, if one received a large amount of positive information about an individual but then received a new piece of negative information, the new negative information would be expectancy violating due to its inconsistency with the available target-specific information. Conversely, if one received a large amount of negative information about an individual but then received a new piece of positive information, the new positive information would be expectancy violating due to its inconsistency with the available target-specific information. Thus, positive and negative information should be equally expectancy violating to the extent that they conflict with a large amount of target-specific information.

In contrast, when little target-specific information is available, perceivers' reference of comparison should be jointly determined by the available target-specific information and perceivers' positive default beliefs about other people in general. In this case, the relative strength of expectancy violation in response to new information depends not only on the (in)consistency of this information with prior target-specific information, but also on its (in)consistency with perceivers' positive default beliefs (e.g., Cacioppo et al., 1997; Peeters & Czapinski, 1990). For example, a new piece of negative information about a target person would be inconsistent with both prior positive information about the target and positive default beliefs, and therefore result in a relatively strong expectancy violation. Conversely, a new piece of positive information about a target person would be inconsistent with prior negative information about the target but consistent with positive default beliefs, and therefore result in a relatively weak expectancy violation. Thus, in the absence of large amounts of target-specific information, negative information should produce greater expectancy violations than positive information, because

positive default beliefs may determine perceivers' reference of comparison over and above the available target-specific information.

THE CURRENT RESEARCH

The goal of the current research was to investigate whether the emergence of a negativity bias in expectancy violation depends on the amount of prior information that is available about a specific target. Based on the assumptions of expectancy-contrast and frequency-weight theories, we hypothesized that new negative information should result in greater expectancy violations than new positive information when prior impressions of a target person are based on relatively little information. In contrast, when prior impressions of a target person are based on large amounts of information, new negative information and new positive information should result in equal expectancy violations.¹

EXPERIMENT 1

To test these hypotheses, participants in Experiment 1 were asked to form an impression of a target individual based on positive or negative behavioral statements that were presented one-by-one against different background colors. Following these statements, participants were presented with a critical target statement that was either congruent or incongruent with the valence of the initial statements. After the impression formation task, participants completed a surprise recognition test, in which they had to identify the background color against which the target statement was presented during the impression formation task. Based on previous research with this paradigm (e.g., Brannon et al., 2017; Gawronski, Ye, Rydell, & De Houwer, 2014; Ye, Tong, Chiu, & Gawronski, 2017), enhanced memory for the background color of valence-incongruent target statements (compared to valence-congruent target statements) was interpreted as an indicator of expectancy violation. Critically, to test whether negativity bias in expectancy violation depends on the amount of prior target-specific information, half of the participants were provided with little initial information about the target, while the other half received a large amount of initial information (see Srull, 1981; Srull, Lichtenstein, & Rothbart, 1985). Thus, we tested whether (1) the memory advantage for the background color of valence-incongruent target statements differs as a function of the valence of

1. For all studies reported in this article, we report all measures, all conditions, and all data exclusions. The data for each study were collected in one shot without intermittent statistical analyses. Based on prior research using the same experimental paradigm, we aimed for a sample size of 320 participants per study (i.e., 40 participants per cell). Based on the effect sizes for expectancy-violation effects in earlier research using the same experimental paradigm (Brannon et al., 2017), a sample size of 320 provides a power of 80% to detect a significant three-way interaction between impression valence, target valence, and amount of prior information with an odds ratio of 0.73 or 1.38 (depending on the direction of the effect). All materials, data, and analysis files are available at <https://osf.io/z9g75/>.

the target statements, and (2) whether any such valence asymmetry depends on the amount of initial information.²

METHOD

Participants and Design. Participants completed a “psychological study on how people form impressions of other people” on Amazon’s Mechanical Turk (MTurk). Eligibility for participation was limited to MTurk workers who (1) lived in the United States, (2) had a HIT approval rate of at least 95% at the time of the study, and (3) had not participated in prior studies from our lab using the same expectancy-violation paradigm. Participants received \$0.50 for their participation. On average, the study took 10 minutes to complete. Of the 341 participants who began the study, 323 participants completed all components of the study and are included in analyses (195 women, 126 men; $M_{\text{age}} = 35.05$ years, $SD_{\text{age}} = 12.17$ years; demographic data missing for 2 participants).³ Participants were randomly assigned to one condition in a 2 (Impression Valence: positive versus negative) \times 2 (Target Valence: positive versus negative) \times 2 (Prior Information: small amount versus large amount) between-subjects design.

Impression Formation Task. Participants were asked to form an impression of an individual (“Bob”) on the basis of statements describing his behavior. Behavioral statements were presented below the individual’s picture for five seconds, with each behavioral statement appearing on successive screens. Each statement was presented against one of ten different background colors, with the target statements always presented against a blue background. The statements were adapted from a list of positive and negative behaviors used by Rydell and Gawronski (2009). Depending on the particular condition, participants first learned either positive or negative information about Bob, and then read a target statement that was either congruent or incongruent with the valence of the initial statements. Participants in the *small amount* condition were presented with five initial statements, then the target statement of interest, followed by four filler statements. Participants in the *large amount* condition were presented with 20 initial statements, then the target statement of interest, followed by four filler statements. All filler statements matched the valence of the initial statements. The backgrounds were evenly distributed throughout the task, including the blue background of the target statement.

2. Different from previous research on person memory, the current experiments investigate memory for incidental contextual features surrounding social information, rather than memory for the social information itself. Thus, while previous person memory research provided valuable insights into the mental organization and application of social information, the current research aims to better understand the extent to which new social information is deemed as (in)consistent with previous information, which is captured by a memory advantage for incidental contexts of expectancy-incongruent information compared to incidental contexts of expectancy-congruent information.

3. Three participants completed the study, but did not submit a request for compensation. Of the 323 participants who completed the study, 20 participants failed an instructional attention check (see Oppenheimer, Meyvis, & Davidenko, 2009) and 5 indicated that they suffered from some form of color blindness. Excluding these participants did not alter the pattern of results, so they are retained in the analyses.

Background Recognition Task. After the impression formation task, participants completed a surprise recognition test. Participants were asked to identify the background color against which each of seven statements was presented during the impression formation task: three statements randomly selected from the initial and filler statements, the target statement, and three more statements from the initial and filler statements. The statements appeared toward the bottom of the screen below a multiple-choice list consisting of a small square of each color used in the impression formation task. Each response option was labeled with a number from 0 to 9. Participants were asked to indicate their response by choosing one of the number options. Past research suggests that expectancy violations enhance attention (Roese & Sherman, 2007), which in turn leads to better memory for the background color in the surprise recognition test (e.g., Brannon et al., 2017; Gawronski et al., 2014; Ye et al., 2017). Thus, correct identification of the background color of the critical target statement served as the dependent variable.

RESULTS AND DISCUSSION

Using binary logistic regression, Impression Valence, Target Valence, and Prior Information were entered into a block-wise model as dummy-coded predictors with each level of interactions entered into a subsequent block. Analyses revealed a significant two-way interaction between Impression Valence and Target Valence, $B = -2.56$, $SE = 0.54$, Wald $Z = 22.54$, $p < .001$, $OR = 0.08$, replicating the expectancy-violation effect obtained in previous research. Specifically, recognition rates for the background color of the target statement were higher when positive target statements followed negative initial statements ($M = .39$, 95% CI [.30, .49]) than when positive target statements followed positive initial statements ($M = .14$, 95% CI [.04, .23]). Conversely, background recognition rates were higher when negative target statements followed positive initial statements ($M = .42$, 95% CI [.32, .52]) than when negative target statements followed negative initial statements ($M = .19$, 95% CI [.09, .28]). Importantly, this two-way interaction was not qualified by a higher-order interaction with Prior Information, $B = 1.32$, $SE = 1.09$, Wald $Z = 1.46$, $p = .227$, $OR = 3.74$ (see Figure 1).

The absence of a significant three-way interaction suggests that the obtained expectancy-violation effect is unqualified by the amount of prior information about the target. However, it does not guarantee that (1) the size of expectancy-violation effects is unaffected by Target Valence and (2) the size of expectancy-violation effects is unqualified by Target Valence within each level of the Prior Information factor. To address these questions, we recoded the Target Valence factor to reflect the congruence of the target statement with the valence of the initial statements. The resulting scores were submitted to separate 2 (Valence Congruence: congruent versus incongruent) \times 2 (Target Valence: positive versus negative) logistic regressions for each condition of the Prior Information factor. The analyses revealed a significant main effect of Valence Congruence in both the small amount condition, $B = -0.97$, $SE = 0.37$, Wald $Z = 6.95$, $p = .008$, $OR = 0.38$, and the large amount condition, $B = -1.58$, $SE = 0.40$, Wald $Z = 16.07$, $p < .001$, $OR = 0.21$. Counter to the hy-

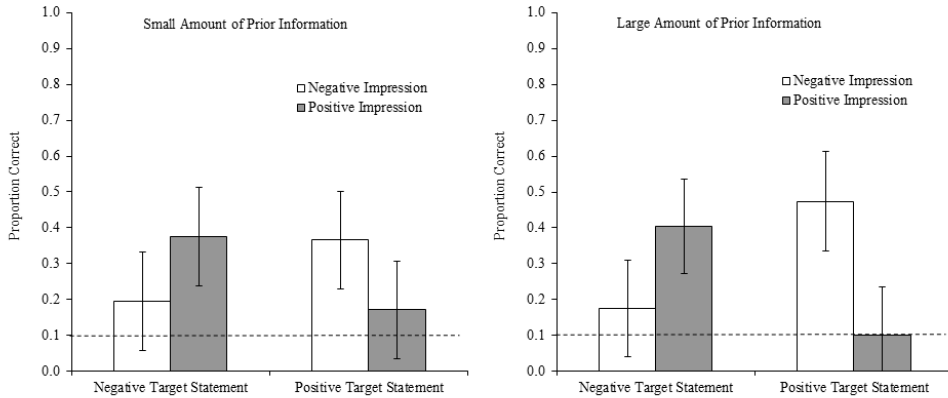


FIGURE 1. Mean proportion of correct background recognition as a function of target statement valence, initial impression valence, and amount of prior information, Experiment 1. Dotted line represents chance responding. Error bars represent 95% confidence intervals.

pothesis that negativity bias in expectancy violation might depend on the amount of prior information about a target, the main effect of Valence Congruence was not qualified by Target Valence in either of the two Prior Information conditions, $B = -0.12$, $SE = 0.73$, Wald $Z = 0.03$, $p = .866$, $OR = 0.88$ and $B = -0.93$, $SE = 0.81$, Wald $Z = 1.31$, $p = .252$, $OR = 0.40$, respectively.

Together, these results suggest that (1) expectancy-violation effects did not differ as a function of whether the expectancy-violating information was negative or positive, and (2) there was no negativity bias in the violation of expectancies regardless of whether participants received a small or large amount of prior information about the target individual.

EXPERIMENT 2

Experiment 1 found no evidence for a negativity bias in expectancy violations, regardless of the amount of prior information about the target. However, a potential limitation of Experiment 1 is that the amount of information participants' received prior to the target statement was confounded with the total length of the task. That is, participants who received little information prior to the target statement were presented with fewer statements overall than participants who received a large amount of information prior to the target statement (i.e., 10 versus 25 state-

4. Although the procedural change in Experiment 2 resolves the confound in Experiment 1, it introduces a new confound between the amount of information participants' received prior to the target statement and the delay between the target statement and the recognition task. Thus, each study contains a unique limitation that is addressed by the respective other. To the extent that the two experiments produce the same pattern of results, the converging evidence suggests that the unique limitation of each study is irrelevant for the obtained results.

ments). Experiment 2 aimed to address this limitation by keeping the total number of statements constant across conditions.⁴

METHOD

As in Experiment 1, participants completed a “psychological study on how people form impressions of other people” on Amazon’s MTurk. Eligibility for participation was limited to MTurk workers who (1) lived in the United States, (2) had a HIT approval rate of at least 95% at the time of the study, and (3) had not participated in prior studies from our lab using the same expectancy-violation paradigm. Participants received \$0.50 in return for their participation. On average, the study took 12 minutes to complete. Of the 364 participants who began the study, 326 participants completed all components of the study and are included in the analyses (170 women, 151 men; $M_{\text{age}} = 34.84$, $SD_{\text{age}} = 12.82$; demographic data missing for 5 participants).⁵ Participants were randomly assigned to one condition in a 2 (Impression Valence: positive versus negative) \times 2 (Target Valence: positive versus negative) \times 2 (Prior Information: small amount versus large amount) between-subjects design.

The procedure and materials were identical to Experiment 1, the only exception being that all participants were presented with 30 statements about the target individual. Participants in the *small amount* condition were presented with five initial statements, then the target statement of interest, followed by 24 filler statements. Participants in the *large amount* condition were presented with 20 initial statements, then the target statement of interest, followed by nine filler statements. The ten backgrounds, including the blue background of the target statement, were evenly distributed throughout the task using the semi-randomized blocked design.

RESULTS AND DISCUSSION

Data were analyzed in line with the procedures of Experiment 1. Analyses revealed a significant two-way interaction between Impression Valence and Prior Information, $B = -1.68$, $SE = 0.62$, Wald $Z = 7.26$, $p = .007$, $OR = 0.19$. When the initial valence was negative, recognition rates for the background color of the target statement were higher when the target statement appeared earlier in the sequence ($M = .28$, 95% CI [.19, .36]) than when it appeared later ($M = .16$, 95% CI [.07, .25]). Conversely, when the initial valence was positive, recognition rates were higher when the target statement appeared later in the sequence ($M = .32$, 95% CI [.23, .41]) than when it appeared earlier ($M = .17$, 95% CI [.08, .26]). Because this effect is

5. Six participants completed the study, but did not submit a request for compensation. Of the 326 participants who completed the study, 9 participants failed an instructional attention check (see Oppenheimer et al., 2009) and 3 participants indicated that they suffered from some form of colorblindness. Excluding these participants did not alter the pattern of results, so all data are retained for analyses.




FIGURE 2. Mean proportion of correct background recognition as a function of target statement valence, initial impression valence, and amount of prior information, Experiment 2. Dotted line represents chance responding. Error bars represent 95% confidence intervals.

independent of the manipulation of Target Valence, it is irrelevant for our concern with expectancy violations.

More important for our main question, analyses also revealed a significant two-way interaction between Impression Valence and Target Valence, $B = -2.98$, $SE = 0.63$, Wald $Z = 22.56$, $p < .001$, $OR = 0.05$. Replicating the findings of Experiment 1, recognition rates for the background color of the target statement were higher when positive target statements followed negative initial statements ($M = .34$, 95% CI [.25, .43]) than when positive target statements followed positive initial statements ($M = .13$, 95% CI [.04, .21]). Conversely, recognition rates were higher when negative target statements followed positive initial statements ($M = .36$, 95% CI [.27, .45]) than when negative target statements followed negative initial statements ($M = .10$, 95% CI [.01, .19]). Importantly, this two-way interaction was not qualified by a higher-order interaction with Prior Information, $B = -0.70$, $SE = 1.25$, Wald $Z = 0.31$, $p = .576$, $OR = 0.50$ (see Figure 2).

As in Experiment 1, we recoded the Target Valence factor to reflect the congruence of the target statement with the valence of the initial statements. Replicating the findings of Experiment 1, separate 2 (Valence Congruence: congruent versus incongruent) \times 2 (Target Valence: positive versus negative) logistic regressions for the two Prior Information conditions revealed a significant main effect of Valence Congruence in both the small amount condition, $B = -1.68$, $SE = 0.44$, Wald $Z = 14.58$, $p < .001$, $OR = 0.19$, and the large amount condition, $B = -1.23$, $SE = 0.40$, Wald $Z = 9.43$, $p = .002$, $OR = 0.29$. Counter to the hypothesis that negativity bias in expectancy violation might be limited to conditions in which little prior information about a target is available, the main effect of Valence Congruence was not qualified by Target Valence when participants received a small amount of prior information, $B = -1.37$, $SE = 0.91$, Wald $Z = 2.29$, $p = .130$, $OR = 0.25$. However, when participants received a large amount of prior information, there was a significant interaction between Valence Congruence and Target Valence, $B = 1.94$, $SE = 0.87$,

Wald $Z = 5.00$, $p = .025$, $OR = 6.93$, indicating a stronger expectancy-violation effect for negative compared to positive target statements.

Together, these results corroborate the conclusions from Experiment 1 that (1) expectancy-violation effects do not differ as a function of whether the expectancy-violating information is negative or positive and (2) there is no negativity bias even when there is little prior information about the target. If anything, the findings of Experiment 2 suggest a negativity bias for expectancy violations when a large amount of target-specific information is available. This finding directly contradicts our hypothesis that negativity bias in expectancy violation might be limited to conditions when a small amount of target-specific information is available. However, because the experimental procedure in this condition was identical to the default condition in previous studies, and these studies failed to obtain any evidence for valence asymmetries in three studies with close to 2000 participants (see Brannon et al., 2017), we refrain from drawing any conclusion from this particular finding.

EXPERIMENT 3

The results from Experiments 1 and 2 disconfirm the hypothesis that new negative information leads to greater expectancy violations than new positive information when the new information conflicts with small amounts of target-specific information. Instead, we found that new negative and new positive information resulted in equal expectancy violations even when the amount of prior information was relatively small. One potential explanation of this finding is that participants formed strong target-specific expectancies on the basis of what we considered a small amount of prior information (i.e., five behavioral statements). This interpretation would be consistent with research on the correspondence bias (for a review, see Gawronski, 2004) and spontaneous trait inference (for a review, see McCarthy & Skowronski, 2014), which suggests that people draw strong inferences about underlying dispositions based on relatively little behavioral information. In line with this interpretation, a follow-up study found that participants reported equally strong expectancies when receiving either 5 or 20 pieces of information (see Supplemental Materials). Thus, to investigate whether the absence of a negativity bias in the small amount conditions of Experiments 1 and 2 was due to the presence of strong target-specific expectancies, Experiment 3 used a larger discrepancy in the manipulation of small versus large amounts of prior information. Specifically, in the condition in which participants received little prior information, only one statement appeared before the target statement. Conversely, in the condition in which participants received a large amount of prior information, 21 statements appeared before the target statement. The effectiveness of this manipulation in producing target-specific expectancies of differential strength was supported in a second follow-up study, showing that participants who read 21 statements about an individual reported stronger expectancies than participants who read only one statement about an individual (see Supplementary Materials).

METHOD

As in Experiments 1 and 2, participants were recruited via MTurk. Participation was restricted to (1) MTurk workers in the United States who (2) had participated in at least one other study on MTurk, (3) had an approval rating of 95%, and (4) had not participated in other studies conducted by our lab using a similar paradigm. Of the 353 participants who initially began the study, 320 participants submitted for payment on MTurk. Three additional participants completed the critical components of the study but did not submit payment requests before the study expired. Thus, a final sample of 323 is included in the analyses (190 women, 133 men; $M_{\text{age}} = 37.54$ years, $SD_{\text{age}} = 12.83$ years).⁶ As in the previous experiments, participants were randomly assigned to one condition in a 2 (Impression Valence: positive versus negative) \times 2 (Target Valence: positive versus negative) \times 2 (Prior Information: small amount versus large amount) between-subjects design.

The procedure of the impression formation task followed the one outlined in Experiment 1, the only exception being the amount of information participants read in each condition. Participants in the *small amount* condition first read one initial statement, then the target statement, followed by eight filler statements. Thus, participants in this condition read ten total statements, which ensured that each background color was presented once. Participants in the *large amount* conditions read 21 initial statements, then the target statement, followed by eight filler statements. Thus participants in this condition read 30 total statements, with each background color appearing three times. The surprise recognition task was identical to the one in Experiments 1 and 2.

RESULTS AND DISCUSSION

As in Experiments 1 and 2, participants' recognition of the background color for the target statement was submitted to a 2 (Impression Valence) \times 2 (Target Valence) \times 2 (Prior Information) binary logistic regression. The analysis revealed significant main effects of Impression Valence, $B = -0.67$, $SE = 0.27$, Wald $Z = 6.33$, $p = .012$, $OR = 0.51$, and Target Valence, $B = 0.70$, $SE = 0.26$, Wald $Z = 7.05$, $p = .008$, $OR = 2.02$. As in Experiments 1 and 2, these main effects were qualified by a significant interaction between Impression Valence and Target Valence, $B = -2.22$, $SE = 0.57$, Wald $Z = 15.25$, $p < .001$, $OR = 0.11$. This interaction indicated that recognition rates for the background color of the target statement were higher when positive target statements followed negative initial statements ($M = .23$, 95% CI [.14, .32]) than when positive target statements followed positive initial statements ($M = .16$, 95% CI [.07, .25]). Conversely, recognition rates were higher when negative target statements followed positive initial statements ($M = .49$, 95% CI [.39, .58]) than when

6. Twelve participants failed an instructional attention check (see Oppenheimer et al., 2009) and 3 participants indicated that they suffered from some form of colorblindness. Excluding these participants did not alter the pattern of results, so all data are retained for analyses.




FIGURE 3. Mean proportion of correct background recognition as a function of target statement valence, initial impression valence, and amount of prior information, Experiment 3. Dotted line represents chance responding. Error bars represent 95% confidence intervals.

negative target statements followed negative initial statements ($M = .15$, 95% CI [.06, .25]). As in the previous two experiments, this two-way interaction was not qualified by a higher-order interaction with Prior Information, $B = 0.75$, $SE = 1.16$, Wald $Z = 0.41$, $p = .521$, $OR = 2.11$. (see Figure 3).

Following the procedures in Experiments 1 and 2, we recoded the Target Valence factor to reflect the congruence of the target statement with the valence of the initial statements, and then conducted separate 2 (Valence Congruence: congruent versus incongruent) \times 2 (Target Valence: positive versus negative) logistic regressions for the two Prior Information conditions. For the small amount condition, there was a significant main effect of Target Valence, $B = 0.90$, $SE = 0.37$, Wald $Z = 5.97$, $p = .015$, $OR = 2.46$. More important for the current investigation, there was also a significant main effect of Valence Congruence in both the small amount condition, $B = -1.00$, $SE = 0.37$, Wald $Z = 7.21$, $p = .007$, $OR = 0.37$, and in the large amount condition, $B = -1.27$, $SE = 0.42$, Wald $Z = 9.24$, $p = .002$, $OR = 0.28$. In the small amount condition, this main effect of Valence Congruence was not qualified by an interaction with Target Valence, $B = -0.48$, $SE = 0.75$, Wald $Z = 0.40$, $p = .526$, $OR = 0.62$, replicating the results of Experiments 1 and 2. However, in the large amount condition, there was a significant interaction between Valence Congruence and Target Valence, $B = -2.15$, $SE = 0.89$, Wald $Z = 5.84$, $p = .016$, $OR = 0.12$, indicating that positive information following an impression of the opposite valence resulted in greater expectancy violations than did negative information.

Together, these results support the conclusions from Experiments 1 and 2 that (1) expectancy-violation effects do not differ as a function of whether the expectancy-violating information is negative or positive and (2) there is no negativity bias even when there is little prior information about the target. Although the results from Experiment 3 suggest a valence asymmetry in the large amount of information condition, we refrain from drawing any conclusions on the basis of this interaction. Aside from the fact that it is counter to our prediction of a stronger negativity bias under conditions of little prior information, the obtained asymmetry involved greater expectancy violations for positive than negative information,

which is (1) opposite to the unexpected asymmetry obtained in Experiment 2 and (2) in contrast to the results of past research which failed to obtain any evidence for valence asymmetries under conditions of large amounts of prior information (Brannon et al., 2017).

GENERAL DISCUSSION

The goal of the current research was to test whether the emergence of a negativity bias in expectancy violation depends on the amount of prior information that is available about a specific target. Based on extant accounts of the negativity bias (for a review, see Skowronski & Carlston, 1989), we predicted that new negative and new positive information would result in equal expectancy violations when the new information conflicts with large amounts of target-specific information. In contrast, new negative information was hypothesized to result in greater expectancy violations than new positive information when the new information conflicts with small amounts of target-specific information. Contrary to the latter hypothesis, the current studies found that new negative and new positive information led to equal expectancy violations regardless of the amount of prior information.

THEORETICAL IMPLICATIONS

In line with the current results, previous research consistently failed to obtain evidence for a negativity bias in expectancy violations (e.g., Brannon et al., 2017; Cacioppo et al., 1993). Together, the available evidence suggests that there are no valence asymmetries in the violation of expectancies. Although this conclusion may seem somewhat surprising, it does not invalidate research on the ubiquity of the negativity bias. Instead, the work on expectancy violation provides deeper insights into the conditions under which the negativity bias is more likely to occur.

In our view, the lack of a negativity bias in expectancy violation is best explained by distinguishing between different stages in the processing of inconsistencies. Gawronski and Brannon (in press) proposed a three-stage model of inconsistency processing in which (1) an inconsistency is first identified which (2) elicits aversive feelings that (3) people aim to reduce by resolving the inconsistency. As discussed by Brannon et al. (2017), past research and theorizing on the negativity bias were primarily concerned with the *integration* of conflicting information in updating impressions. For example, in a study by Reeder and Coovert (1986), participants updated their initial impressions of an individual's moral character in response to impression-inconsistent information when the new information was negative, but not when it was positive. These results provide clear evidence for a negativity bias, but the focus of the study was on *updating* impressions (i.e., reconciling an inconsistency after it had been identified). Further, the theories reviewed above explicitly apply to how information is integrated into an overall impression after acquiring new information about an individual that is inconsistent in valence. For

example, some theories suggest that immoral behavior that conflicts with an initial impression of a person as being moral leads to a dispositional attribution, whereas moral behavior that conflicts with an initial impression of a person as being immoral leads to a situational attribution (e.g., Reeder & Brewer, 1979). Although the two ways of attributing impression-inconsistent behavior have different implications for belief updating (i.e., updating in response to new information about immoral behavior, but not moral behavior), either one aims to reconcile the inconsistency between the initial impression and the new information. In this sense, the processes proposed by extant accounts of the negativity bias fall into Gawronski and Brannon's (in press) third stage of inconsistency processing (i.e., the reconciliation of inconsistent information).

This emphasis differs from the one in the current studies and earlier research on expectancy violation (e.g., Brannon et al., 2017; Cacioppo et al., 1993), which focus on the first stage of inconsistency processing: the *identification* of inconsistency. The paradigms employed in this work aim to identify the conditions under which new information elicits a surprise response on the basis of prior information (see Noordewier, Topolinski, & Van Dijk, 2016). Thus, their outcomes reflect the perception of inconsistency between the new information and prior information. This conclusion is consistent with findings by Brannon et al. (2017, Experiment 4) who explicitly asked participants whether they viewed a specific piece of information as inconsistent with the information they previously learned about an individual. This approach also failed to produce any evidence for a negativity bias, in line with the results of the memory-based measure of surprise employed in the current work (see also Brannon et al., 2017, Experiments 1–3). Together, these results provide converging evidence that there is no negativity bias in the identification of inconsistency. Instead, the negativity bias seems to arise from different ways of reconciling inconsistency between initial impressions and new information after inconsistency has been identified.

This conclusion highlights the importance of distinguishing between stages of inconsistency processing, as it is possible for effects to differ across these stages. Thus, in addition to providing valuable insights into a potential boundary condition of the negativity bias, the current findings have important implications for the broader work on cognitive consistency, in that it provides a basis for studying potential differences across processing stages. Future research distinguishing between different stages of inconsistency processing may lead to refined theoretical accounts of cognitive (in)consistency, its impact on information processing, and the implications of these processes for the negativity bias in the processing of conflicting information.

POTENTIAL OBJECTIONS

A potential objection to our proposed explanation is that the lack of a negativity bias in expectancy violations could be explained by averaging models of information integration (e.g., Anderson, 1971) without invoking any assumptions re-

garding stages of information processing. Averaging models suggest that overall impressions are formed by taking the average of all available information about an individual and assigning the trait or evaluation that corresponds to that average (e.g., Anderson & Birnbaum, 1976; Yamagishi & Hill, 1981). To illustrate this process, imagine a scale from +5 to -5 where the numeric value indicates the extremity of a given piece of information and the sign indicates the valence of that information. If someone receives one piece of positive information and one piece of negative information that are equal in terms of their extremity (e.g., +2 and -2), the overall evaluation of that individual should be neutral (i.e., 0). In contrast, if the positive piece of information is less extreme than the negative piece of information (e.g., +2 and -4), the overall evaluation of the individual should be negative (i.e., -2).

According to averaging models, positive and negative information that is inconsistent with prior information may produce equal expectancy violations even when negative information has a greater impact on overall impressions than positive information.⁷ To illustrate the difference between expectancy violation and information integration, suppose that positive information, on average, is valued at +2 and negative information, on average, is valued at -4. In this case, negative information would have greater impact on overall impressions due its greater extremity. Nevertheless, positive and negative information should produce the same expectancy violation, because the net difference between the two is the same regardless of whether new negative information conflicts with prior positive information or new positive information conflicts with prior negative information. In the above example, new negative information (-4) that conflicts with prior positive information (+2) would result in a net difference of 6. Similarly, new positive information (+2) that conflicts with prior negative information (-4) would also result in a net difference of 6. Thus, even if negative information is, on average, more extreme than positive information (Alves et al., 2017; Vogt, De Houwer, Koster, Van Damme, & Crombez, 2007), the expectancy violation produced in each case would be equivalent because the net difference would be the same.

Although these arguments explain why the negativity bias in overall impressions may not generalize to expectancy violations, such an account is based on the premise that negativity bias in overall impressions is due to differences in the extremity of positive and negative information. However, several theories of dispositional inference suggest that the negativity bias in overall impressions can be entirely independent of such differences in extremity (see Reeder & Brewer, 1979; Skowronski & Carlston, 1989). According to these theories, negative information can have a greater impact on overall impressions than positive information even when the two pieces of information are equal in terms of their extremity. The reason for this asymmetry is that expectancy-violating negative information is typically attributed to dispositional factors, whereas expectancy-violating positive information tends to be attributed to situational factors. For example, immoral behavior that conflicts with prior moral behavior of the same extremity is attributed

7. We thank Hans Alves for this insight.

to dispositional factors, whereas moral behavior that conflicts with prior immoral behavior of the same extremity is attributed to situational factors (see Reeder & Brewer, 1979). As a result, expectancy-violating immoral behavior leads to greater impression updating than expectancy-violating moral behavior (e.g., Reeder & Coovert, 1986). Thus, averaging models of information integration can explain the absence of a negativity bias in expectancy violation in cases when differences in the extremity of negative and positive information produce a negativity bias in information integration. However, they are unable to explain instances of the negativity bias in information integration that do not involve any differences in terms of extremity.

Another potential concern is that participants in the current studies might have been insufficiently involved in the impression formation task to experience an affective response to inconsistency. Some theories suggest that affective responses to inconsistency depend on strong personal involvement and high subjective importance. For example, inconsistencies may be experienced as aversive only when they threaten a cherished belief or one's self-concept (e.g., Aronson, 1968, 1969; Festinger, 1957; Steele & Liu, 1983). From these perspectives, it seems possible that the current studies failed to obtain evidence for a negativity bias in expectancy violations, because participants were not sufficiently involved in the impression formation task. However, there are at least two reasons why such an explanation is unlikely to account for the current findings. First, some newer research and theoretical perspectives pose a challenge to the notion that personal involvement is necessary for the elicitation of inconsistency effects. For example, people have been found to compensate for inconsistencies as unimportant as anomalous playing cards (e.g., Slegers, Proulx, & van Beest, 2015). Based on these and various other findings, Proulx, Inzlicht, and Harmon-Jones (2012) argued that violations of expectancies elicit aversive feelings regardless of personal involvement and subjective importance. Second, although there was no valence asymmetry in the current research, we did find strong expectancy-violation effects in line with past research (e.g., Brannon et al., 2017). If low involvement undermined effects of inconsistency, there should have been no memory advantage for the background of expectancy-incongruent information. Yet, participants in all three studies showed better memory for the background color of the target statement when the target statement was incongruent than when it was congruent with the valence of their initial impression. Together, these results suggest that lack of personal involvement or low subjective importance do not explain the current pattern of results.

CONCLUSION

Although the tendency for negative information to have a greater impact than positive information is a well-established psychological phenomenon, little work has investigated its boundary conditions and its occurrence at different stages of inconsistency processing. The current work tested whether the emergence of a negativity bias in the violation of expectancies depends on the amount of target-

specific information, but did not find any evidence supporting this moderation. Together with earlier research, these results suggest that the negativity bias in impression formation arises after inconsistency has been identified, not during the identification of inconsistency.

SUPPLEMENTARY MATERIALS

SUPPLEMENTARY EXPERIMENT S1

Experiments 1 and 2 reported in the main text suggest that positive and negative information produce equal expectancy violations, regardless of the amount of target-specific information perceivers received about an individual. However, it is possible that participants formed strong target-specific expectancies on the basis of what we considered a small amount of prior information. If this was the case, the lack of moderation by amount of information could simply be due to the fact that participants in the small amount of information condition used their target-specific expectancies as a reference of comparison for new information, rather than having to rely on their expectations regarding people in general. Supplementary Experiment S1 aimed to test this possibility by examining expectancy strength as a function of 5 versus 20 pieces of information.

Method. Participant recruitment followed the same procedures outlined in the main text. Of the 340 participants who initially began the study, 320 submitted requests for payment on MTurk. Five additional participants completed the study but did not submit for payment prior to the study's expiration. Thus, 325 participants completed all critical parts of the study and are included in the analyses (170 women, 152 men; $M_{\text{age}} = 37.32$, $SD_{\text{age}} = 12.00$; demographic data missing for 3 participants). Participants were randomly assigned to the four conditions of a 2 (Impression Valence: positive versus negative) \times 2 (Amount of Information: small amount versus large amount) between-subjects design. Participants in the small amount condition received 5 statements about a target individual; participants in the large amount condition received 20 statements about a target individual. The valence of the statements was either positive or negative. In all conditions, each statement was presented below the targets picture for five seconds against a white background. Following the last statement, participants completed a measure of expectancy strength which was comprised of six items presented in a random order (see Table 1).

Results. Participants' responses were coded such that higher numbers represent stronger positive expectancies. A composite score was then created by averaging responses to all six items (Cronbach's $\alpha = .98$). Participants' composite expectancy score was then submitted to a 2 (Impression Valence: positive versus negative) \times 2 (Amount of Information: small amount versus large amount) ANOVA. This analysis revealed a significant main effect of Impression Valence $F(1, 322) = 2281.16$, $p < .001$, $\eta_p^2 = .88$. This main effect suggests that participants reported more positive expectancies after receiving positive information about the target ($M = 6.20$, $SD = 0.94$) than after receiving negative information about the target ($M = 1.57$, $SD =$

TABLE 1. Items Assessing Expectancy Strength, Supplementary Experiments 1 and 2

| Item | Scale |
|---|--|
| 1. How likely do you think it is that Bob will behave positively in the future? | 1 (Extremely unlikely) to 7 (Extremely likely) |
| 2. How likely do you think it is that Bob will behave negatively in the future?* | 1 (Extremely unlikely) to 7 (Extremely likely) |
| 3. How confident are you that Bob will behave positively in the future? | 1 (Not at all confident) to 7 (Very confident) |
| 4. How confident are you that Bob will behave negatively in the future?* | 1 (Not at all confident) to 7 (Very confident) |
| 5. Please rate your agreement with the following statement: "I expect Bob to behave positively in the future." | 1 (Strongly disagree) to 7 (Strongly agree) |
| 6. Please rate your agreement with the following statement: "I expect Bob to behave negatively in the future."* | 1 (Strongly disagree) to 7 (Strongly agree) |

Note.*indicates reverse scored item.

0.82). However, this main effect was not qualified by Amount of Information, $F(1, 322) = 1.68, p = .20, \eta_p^2 = .01$, suggesting that participants formed equally strong expectancies after receiving either 5 or 20 pieces of information.

SUPPLEMENTARY EXPERIMENT S2

The results of Supplementary Experiment S1 suggest that participants formed equally strong expectancies on the basis of 5 and 20 pieces of information. Thus, it is possible that participants in both Prior Information conditions relied on their target-specific expectancies as a reference for evaluating new information. To rule out this possibility, we increased the discrepancy between the small and large amount of information conditions, such that participants in the small amount of information condition formed weaker target-specific expectancies. Supplementary Experiment S2 investigated whether a single piece of information leads to weaker expectancies than 20 pieces of target-specific information.

Method. Participant recruitment followed the same procedures outlined in the main text. Of the 338 participants who initially began the study, 320 submitted requests for payment on MTurk. Three additional participants completed the study but did not submit for payment prior to the study’s expiration. Thus, 323 participants completed all critical parts of the study and are included in the analyses (180 women, 140 men; $M_{age} = 36.14, SD_{age} = 11.55$; demographic data missing for 3 participants). As in Supplementary Experiment S1, participants were randomly assigned to the four conditions of a 2 (Impression Valence: positive versus negative) \times 2 (Amount of Information: small amount versus large amount) between-subjects design. The methods were identical to those in Supplementary Experiment S1, except that participants in the small amount of information condition received only a single piece of information.

Results. As in Supplementary Experiment S1, participants' responses were recoded such that higher numbers represented stronger positive expectancies and were averaged to form a single composite score (Cronbach's $\alpha = .97$). This composite score was then submitted to a 2 (Impression Valence: positive versus negative) \times 2 (Amount of Information: small amount versus large amount) ANOVA. The analysis revealed a significant main effect of Impression Valence, $F(1, 319) = 625.32, p < .001, \eta_p^2 = .66$, which was qualified by a significant interaction with Amount of Information, $F(1, 319) = 165.32, p < .001, \eta_p^2 = .34$. To decompose this interaction, separate t -tests comparing the effect of Amount of Information at each level of Impression Valence were constructed. These comparisons revealed that participants formed stronger valence-congruent expectancies in both the positive condition, $t(155) = -7.95, p < .001$, Hedges' $g = 1.27$, and in the negative condition, $t(164) = 10.69, p < .001$, Hedges' $g = -1.66$.

REFERENCES

- Alves, H., Koch, A. S., & Unkelbach, C. (2017). Why good is more alike than bad: Processing implications. *Trends in Cognitive Sciences, 21*, 72-82.
- Anderson, N. H. (1971). Integration theory and attitude change. *Psychological Review, 78*, 171-206.
- Anderson, N. H., & Birnbaum, M. H. (1976). Test of an additive model of social influence. *Journal of Personality and Social Psychology, 33*, 655-662.
- 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.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5*, 323-370.
- Birnbaum, M. H. (1972). Morality judgments: Tests of an averaging model. *Journal of Experimental Psychology, 93*, 35-42.
- 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.
- Cacioppo, J. T., Crites, S. L., Jr., Berntson, G. G., & Coles, M. G. H. (1993). If attitudes affect how stimuli are processed, should they not affect the event-related brain potential? *Psychological Science, 4*, 108-112.
- Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1997). Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review, 1*, 3-25.
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology, 87*, 293-311.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row Peterson.
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology, 38*, 889-906.
- Gawronski, B. (2004). Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias. *European Review of Social Psychology, 15*, 183-217.
- Gawronski, B., & 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 ed.). Washington, DC: American Psychological Association.
- 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.
- Helson, H. (1964). *Adaptation-level theory*. New York: Harper.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in social psychology. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219-266), New York: Academic.
- Kanouse, D. E., & Hanson, L. R., Jr. (1972). Negativity in evaluations. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 47-62). Hillsdale, NJ: Erlbaum.
- McCarthy, R. J., & Skowronski, J. J. (2014). Controlled processing and automatic processing in the formation of spontaneous trait inferences. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 221-234). New York: Guilford.
- Noordewier, M. K., Topolinski, S., & Van Dijk, E. (2016). The temporal dynamics of surprise. *Social and Personality Psychology Compass, 10*, 136-149.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology, 45*, 867-872.
- Peeters, G., & Czapinski, J. (1990). Positive-negative asymmetry in evaluations: The distinction between affective and informational negativity effects. *European Review of Social Psychology, 1*, 33-60.
- Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology, 61*, 380-391.
- 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.
- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review, 86*, 61-79.
- Reeder, G. D., & Coovert, M. D. (1986). Revising an impression of morality. *Social Cognition, 4*, 1-17.
- Roese, N. J., & Sherman, J. W. (2007). Expectancies. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (2nd Ed., pp. 91-115). New York: Guilford.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review, 5*, 296-320.
- 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.
- Sherif, M., & Sherif, C. W. (1967). Attitudes as the individual's own categories: The social judgment approach to attitude change. In C. W. Sherif & M. Sherif (Eds.), *Attitude, ego involvement, and change* (pp. 105-139). New York: Wiley.
- Shook, N. J., Fazio, R. H., & Eiser, J. R. (2007). Attitude generalization: Similarity, valence, and extremity. *Journal of Experimental Social Psychology, 43*, 641-647.
- Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin, 105*, 131-142.
- 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.
- Srull, T. K. (1981). Person memory: Some tests of associated storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 440-463.
- Srull, T. K., Lichtenstein, M., & Rothbart, M. (1985). Associate storage and retrieval processes in person memory. *Journal of Experimental Psychology: Human Learning and Memory, 11*, 316-345.

- Steele, C. M., & Liu, T. J. (1983). Dissonance processes as self-affirmation. *Journal of Personality and Social Psychology, 45*, 5-19.
- Vogt, J., De Houwer, J., Koster, E. H. W., Van Damme, S., & Crombez, G. (2007). Allocation of spatial attention to emotional stimuli depends upon arousal and not valence. *Emotion, 8*, 880-885.
- Wyer, R. S., Jr. (1973). Category ratings as "subjective expected values": Implications for attitude formation and change. *Psychological Review, 80*, 446-467.
- Yamagishi, T., & Hill, C. T. (1981). Adding versus averaging models revisited: A test of a path-analytic integration model. *Journal of Personality and Social Psychology, 41*, 13-25.
- 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.