

Individual Differences in Valence Weighting: When, How, and Why They Matter

Matthew D. Rocklage^a and Russell H. Fazio^{a*}

^aOhio State University, 1835 Neil Avenue, Columbus, OH 43210, USA

Journal of Experimental Social Psychology, in press

* Corresponding author at: Department of Psychology, Ohio State University, 1835 Neil Avenue, Columbus, OH 43210, USA. Fax: +1 614 292 5601

E-mail addresses: rocklage.1@osu.edu (M. D. Rocklage), fazio.11@osu.edu (R.H. Fazio)

Abstract

Recent research has shown that individuals vary in the extent to which they weight positive versus negative information during attitude generalization, i.e., their valence *weighting bias* (Pietri, Fazio, & Shook, 2013). As of yet, little is known about the conditions under which such valence weighting is likely to affect behavior and the consequences of that behavior.

Experiments 1 and 2 tested the idea that the relative weight individuals give to positives versus negatives may influence their formation of an initial evaluative response, which will serve as a default provided that they do not have the motivation and opportunity to deliberate further. When opportunity was restricted by the requirement to respond quickly, participants showed greater correspondence between their weighting bias and their approach-avoidance behavior toward objects in a novel environment (Experiment 1). When an experimental manipulation motivated them to mistrust their initial responses, participants showed less correspondence between their weighting bias and risk-taking behavior than when they were motivated to trust their initial responses (Experiment 2). Experiment 3 investigated the downstream consequences of this valence weighting bias for attitude maintenance versus change. Those with a more negative weighting bias gave greater weight to negative information that was actually false, avoided testing the associated stimuli, and hence did not discover their true value. Those with a more positive weighting bias gave less weight to the negative information, tested the associated stimuli more fully, and overcame the false negative information. Implications for exploration, attitude maintenance, and prejudice are discussed.

Keywords: attitudes, valence weighting, negativity bias, exploration, risk, stereotypes

Individual Differences in Valence Weighting: When, How, and Why They Matter

Valence is one of the most far-reaching constructs in psychology. It applies to domains as diverse as judgment and decision-making, close relationships, the self, and stereotyping. Not only is valence far-reaching, but much of how we interact with our world is determined by the positive and negative associations we have. For instance, given a positive attitude toward objects or people, we will be all the more likely to engage with them, thereby leading to the potential experience of positive outcomes, as well as information about the validity of the positive association. A negative attitude, on the other hand, can lead to avoidance behavior, thereby evading potential harm, but also forgoing any possible benefits of interaction with the object or person. Such avoidance behavior can maintain our original negative association without actually testing its validity (Fazio, Eiser, & Shook, 2004).

When investigating valence, researchers have often put forward the possibility that individuals tend to be affected more by negatives than positives (e.g., Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Rozin & Royzman, 2001). Similarly, other researchers have argued that a single-unit increase of negativity has greater implications for subsequent behavior compared to a single-unit increase of positivity (Cacioppo, Gardner, & Bernstein, 1997). Beyond this general asymmetry, researchers have also proposed that individuals differ in the extent to which they focus on positives versus negatives. For example, the behavioral inhibition and activation systems (BIS/BAS) have been put forth to account for individual differences in sensitivity to punishment (BIS) and to reward (BAS; Gray, 1987). Likewise, research on regulatory focus shows that individuals can differ in the extent to which they are sensitive to gains – promotion focused – or sensitive to losses – prevention focused (Higgins, 1997). More recently, approach/avoidance temperament has been proposed as basic dimensions of personality

that capture individual differences in sensitivity to positives and negatives (Elliot & Thrash, 2010). It is these sensitivities, in turn, that are argued to underlie observed relations among different approach-related (e.g., extraversion, positive emotionality, BAS, and promotion focus) and avoidance-related constructs.

Researchers have also recently developed and demonstrated the utility of a performance-based measure of one particular aspect of differential valence sensitivity: the extent to which individuals weight positive versus negative information when generalizing their existing attitudes to novel objects (Pietri, Fazio, and Shook, 2013). The measure stems from a computer game called BeanFest in which participants learn about beans, varying in shape and number of speckles, that produce either positive or negative outcomes when selected. After learning about these beans during a game phase, participants then classify new beans that vary in resemblance to these game beans as being either positive or negative. It is these classifications that provide the basis for the valence weighting measure. For example, when faced with a novel bean that resembles both a positive and negative game bean, an individual who categorizes that bean as negative is weighting negative information more heavily. Indeed, some individuals show evidence of having generalized negative attitudes to a greater extent than positive attitudes, indicating they have given greater weight to negative resemblances. Other individuals generalize their positive attitudes more strongly, weighting resemblance to a known positive more heavily than resemblance to a known negative, thus leading to more favorable appraisals of the novel beans.

Utilizing this behavioral measure of individuals' valence *weighting bias*, Pietri et al. (2013) found that the bias related to judgments of novel events in a wide variety of domains, including interpersonal relationships, threat assessment, and risk propensity. Specifically, a more

negative weighting bias (i.e., generalizing negative attitudes to a greater extent than positive attitudes and therefore giving more weight to negative features than positive features) was related to greater expressed concern about specific situations that allowed for the possibility of social rejection, perceptions that various potentially threatening events were likely to increase in severity, a general apprehension about meeting new people and entering new situations, a decreased propensity to endorse risky options, and more cautious behavior in a gambling situation. Given its relevance across these diverse domains, the weighting of positive versus negative appears to be a fundamental bias that generally characterizes individuals' judgments of novel objects or events. These relations make sense as any such novel judgments are essentially exercises in attitude generalization.

Apart from its fundamental nature, as of yet, relatively little is known about the characteristics of the weighting bias. There are, however, two important findings to note. First, as intimated above, it does appear that the weighting bias is most impactful in novel situations. Indeed Pietri et al. (2013) found that the weighting bias was strongly related to judgments regarding situations that college students were unlikely to have experienced (e.g., chasing tornados to take dramatic photos), while largely unrelated to those they were likely to encounter in the past (e.g., exposing oneself to the sun without sunscreen). It seems likely that once having experienced the situation, individuals can utilize what they learned from that experience and need not rely on valence weighting. Second, individuals appear to have difficulty reporting their valence weighting tendencies. For instance, when directly asked about the extent to which they weight positive versus negative information, individuals' self-reports did not correlate with their weighting tendencies as measured in BeanFest. Similarly, Pietri et al. found that the weighting bias does not overlap with other self-report measures such as

approach/avoidance temperament (Elliot & Thrash, 2010). As has been noted by many researchers, valence is often confounded with distinctiveness and diagnosticity (e.g., Skowronski & Carlston, 1989). Therefore, discerning and reporting their valence weighting tendencies may be particularly challenging for individuals. Measuring individuals' weighting tendencies via the attitude generalization task of BeanFest, on the other hand, is done behaviorally and utilizes novel stimuli that are experimentally associated with positive and negative outcomes and, hence, is free of any such confounds.

Thus although valence weighting tendencies have been established as a fundamental individual difference, when they are most likely to operate and have their largest impact remains largely unknown. The current research therefore seeks to address the following questions: Under what conditions is the weighting bias most likely to operate? Under what conditions is it most likely to have its largest impact? In addition, under what conditions might individuals deviate from their typical valence weighting tendencies? And, finally, what are the downstream consequences of acting on the basis of one's weighting bias?

Knowing when such valence weighing tendencies are most likely to operate provides not only conceptual clarity to the construct itself, but also a better understanding of when, how, and why such tendencies may prove beneficial to individuals. For example, it may be the case that this valence weighting bias allows for efficient decision processes that require relatively few resources, and therefore is especially pivotal when individuals need to make decisions under time pressure or other challenging circumstances. Once the conditions associated with the operation of valence weighting proclivities have been identified, it would then be important to demonstrate what the consequences of relying on such valence weighting might be.

Valence Weighting in the Formation of an Initial Evaluative Response

Given that the process of distilling and integrating positive and negative features appears to occur across a number of domains, it seems that individuals are likely to become quite practiced at such valence weighting. As a result, individuals' valence weighting biases may facilitate their quick appraisal of a novel stimulus and the development of an initial attitude toward the stimulus. Under certain circumstances, this evaluative response may prove sufficient for behavioral decisions. That is, the initial appraisal resulting from the weighting bias may provide an acceptable default basis for action toward the object. One way of illuminating such a possibility is to consider the circumstances under which individuals' weighting proclivities might prove influential from the perspective of dual-process models.

Though there are many different flavors, dual-process models all have similar features. They all postulate a more spontaneous processing method that is often referred to as more automatic, intuitive, top-down, or "quick-and-dirty" (see Chaiken & Trope, 1999) as well as a second, more deliberative processing method, which has been described as more controlled, thoughtful, effortful, or bottom-up. The Motivation and Oppportunity as DEterminants (MODE) model is one such model and is particularly relevant as its very focus concerns attitudes and the multiple processes by which they influence behavior (Fazio, 1990; Olson & Fazio, 2009).

As a dual-process model, the MODE model shares the characteristic postulate regarding spontaneous versus deliberative processing. It also argues for the possibility of "mixed processes" that involve a combination of automatic and controlled components. The spontaneous mode refers to an attitude-to-behavior process in which judgments or behaviors are a downstream consequence of an automatically-activated attitude. Once activated, the attitude biases construal of the object in the immediate situation. These immediate perceptions will determine behavior unless the individual engages in more deliberative processing that serves to

override the initial response. Whether such deliberation is pursued then depends on individuals' *motivation* to engage in more effortful reflection and their *opportunity* to do so (e.g., sufficient time, cognitive resources, or general ability). In other words, only when individuals simultaneously have the motivation and the opportunity to counter the influence of the automatically-activated attitude will they succeed in overriding the immediate perceptions produced by the attitude.

There are many situations, however, when individuals do not have pre-existing attitudes toward the object in question. According to the MODE model, in such situations individuals have no choice but to construct an evaluative judgment on the spot. This construction can at times be very deliberative, but it can also occur through a cursory consideration of the alternatives at hand in order to come to a quick decision. In other words, this construction process can range from more to less extensive depending on the motivation and opportunity factors mentioned previously. Indeed, consider how such a construction process might proceed. It would prove functional for individuals to quickly distill the positive and negative features that characterize the object in question, i.e., its resemblance to known positives versus known negatives. Importantly, these resemblances could range from being more concrete in nature – a novel bean could be visually similar to a previously-encountered bean known to be positive or negative – to more abstract – a situation resembles one that has proven negative in the past, but it also bears some similarity to one that has proven positive. Individuals then could weight these valences against one another to arrive at an initial response. Individuals' general valence weighting tendencies are likely to influence such initial assessments. In that sense, the weighting bias may serve as the basis for the construction of an initial evaluative response. Whether this

suffices as a default basis for action is likely to depend on the individual's motivation and opportunity to engage in further deliberation.

The aim of our first two experiments, then, was to examine the conditions under which valence weighting is most likely to matter from the perspective of the MODE model. If it is the case that valence weighting proclivities influence the formation of initial default responses, their relation to behavior should be most apparent under conditions characterized by little opportunity (Experiment 1) or little motivation to deliberate further (Experiment 2).

Valence Weighting and its Consequences

While Experiments 1 and 2 of this paper are primarily concerned with identifying the conditions under which the weighting bias is most likely to prove influential – under little time and with little motivation to do otherwise – Experiment 3 investigates the downstream consequences of reliance on the weighting bias. As indicated earlier, individuals will be forced to construct attitudes in situations where they have no pre-existing attitudes. It thus seems likely that the weighting bias could have a particularly large impact on the formation and/or maintenance of attitudes in these novel situations. In particular, the weighting bias could have its impact in at least two related ways.

First, the weighting bias could affect the starting point in the construction of an evaluative response. In other words, given initial positive or negative information, those with a more negative weighting bias would likely begin with a more negative attitude than those with a positive weighting bias. Second, the effect of this initial weighting could be seen in its downstream consequences when interacting with those same stimuli subsequently. Specifically, having weighted the initial negative information more heavily, individuals with a more negative weighting bias would be more likely to avoid the stimuli associated with this negative

information and thus maintain their original negative evaluations. As demonstrated by Fazio et al. (2004), avoiding supposedly negative stimuli leads to the maintenance of the negative evaluations even if the negative information is false. Given our current reasoning, this should be particularly true of those with a negative weighting bias. Those with a more positive weighting bias, on the other hand, should be more likely to overcome any initial negative information. They should have given such information less weight and may therefore be more likely to subsequently test its validity by engaging with the stimuli. To explore these possibilities, Experiment 3 investigated the downstream consequences of having a given weighting bias and its effect on attitude change or maintenance.

Taken together, the three current experiments aim to illuminate the conditions under which the valence weighting bias is likely to matter and its consequences, focusing both on our characterization of it as a means toward an initial default response and on the effects it can have on subsequent approach-avoidance behavior and attitude change versus maintenance.

Experiment 1

Utilizing the perspectives provided by the MODE model, we investigated the characteristics of the weighting bias and, in particular, when it would prove influential. In this first experiment, we hypothesized that if the weighting bias is appropriately characterized as influencing an initial default response, then participants should rely on this response to a greater extent when they have little time to do otherwise. To that end, the following experiment restricted the time participants had to make decisions so as to test whether such reduced opportunity to deliberate would increase the correspondence between their weighting bias and their behavior. Furthermore, following past research (Pietri et al., 2013), we also anticipated that this correspondence may be all the more likely when the situation at hand was still novel.

Method

Overview. Upon entering the lab, participants played BeanFest, a computer game which measured their weighting bias on a continuum from relatively positive to negative (Pietri et al., 2013). Participants then played a game similar to BeanFest called DonutFest that involved weighting the potential positive outcome of earning points and gaining information about the novel DonutFest environment against the relative negative outcome of selecting a potentially harmful donut and losing points in the game. Thus, approaching novel stimuli in this paradigm is a risk with potential benefits, but also potential drawbacks. All participants were exposed to BeanFest in the same way, given that it was used to assess their valence weighting tendencies. However, when continuing on to DonutFest, half of the participants had 5000 ms to make a decision of whether to approach or avoid each of the novel donut stimuli (unrestricted time condition)¹ whereas the other half were given only 1000 ms to decide (restricted time condition).

Participants. Participants were 64 (28 male and 36 female) introductory psychology students who completed the experiment for class credit. Two participants were excluded for an extreme number of missed trials (greater than 2 standard deviations above the mean) in the DonutFest paradigm (in other words, a failure to meet the task requirement of responding within the time deadline). Another participant in the unrestricted time condition was excluded for not taking the DonutFest game seriously, (approaching on 108 out of the 120 possible trials (three standard deviations above the mean), failing to respond on 11 trials, and providing only a single avoidance response during the entire game). A total of 61 participants remained for subsequent analyses (27 male and 34 female).

¹ Although time was restricted to 5000 ms, participants rarely failed to respond within this time limit and therefore this condition may be considered functionally similar to a condition with no time constraints.

Procedure.

BeanFest: Assessing individuals' weighting bias. The procedure for BeanFest in the three experiments reported in this paper follows Pietri et al. (2013). Participants began the experiment by reading instructions related to the BeanFest game. They then had six practice trials where they could become comfortable with the buttons and screen displays that would be used throughout the game.

After this practice phase, participants continued on to the first of three game blocks. During each game block, participants were shown 36 out of the total of 100 beans. The beans varied systematically on two dimensions: their shape (10 levels from circular to oblong) and how many speckles they had (1 to 10). The 36 beans shown in each game-phase block were from six regions of this 10x10 matrix (see Pietri et al., 2013, for details regarding the matrix). They were presented in random order and each retained its original value across the three blocks. It was the goal of the participant to learn which beans were positive (+10 points) and which were negative (-10 points). Participants were told they must either choose to approach and select the bean by pressing 'd' on the keyboard, or avoid the bean by pressing 'k.' When selecting a positive bean, participants were told the value of the bean, earned 10 points, and were shown an increase in overall points on their points meter in the lower left side of the screen. Conversely, selecting a negative bean would decrease participants' points by 10. If participants decided to avoid the bean, they would not gain or lose any points, *but would still be told the value of the bean.* In other words, learning a bean's value was not contingent on approaching the bean; participants received information about each bean on each and every trial.²

² BeanFest is sometimes implemented with feedback contingent on approach behavior. Given our present interest in the weighting bias, however, we utilized the full-feedback approach, just as Pietri et al. (2013) did. This

Participants started the game with 50 points, which was intended to provide a neutral stance with respect to gain versus loss framing. The total points a participant could have ranged from 0 points – participants were notified they had lost the game if this occurred – to 100 points – participants were notified they had won the game if this occurred. After winning or losing a game, participants' point total was reset to 50 and they immediately began another game. The beans did not change value after participants won or lost and all participants still saw the same 36 beans three times regardless of how many times they won or lost. These procedural settings follow from past research using BeanFest; for more details, see Fazio et al. (2004).

After the BeanFest game, participants continued on to the test phase in which they were shown all 100 beans in the matrix. In addition to the 36 beans from the game phase, there were also 64 novel beans that varied in resemblance to these previously-seen positive and negative beans. Participants were instructed to categorize the beans as either helpful or harmful and received no feedback on the correctness of their decisions. Points were not awarded and the points meter that had been present in the lower left of the screen during the game phase was no longer visible. Thus, the test phase focused upon how well participants had learned the game beans and the extent to which their positive and negative attitudes generalized to novel beans. As detailed in the Results section, the calculation of participants' weighting bias centers on their responses to these novel beans.

DonutFest: The experimental manipulation and dependent variable. After completing the test phase of BeanFest, participants were told they would be starting a new game involving novel stimuli that resembled donuts. These novel stimuli varied from how yellow to orange to

implementation has the advantage of not producing overall mean differences in the learning of positive vs. negative beans (see Fazio et al., 2004).

red they were in color (10 different levels), and how large the hole in the middle of the stimulus was (10 levels). There were again six practice trials followed by three game-phase blocks. It was anticipated that the three game-phase blocks would provide participants time to learn the valence of the various types of donuts and thereby allow us to also test the possibility that the weighting bias would be most impactful in novel situations, i.e., before participants learned the true value of the various donuts.

The parameters for DonutFest were the same as in BeanFest except that instead of 36 stimuli presented in each game phase, there were now 40. More importantly, in DonutFest participants would not learn about a donut's value unless they chose to approach it. In other words, learning about the donut was contingent on approach behavior. Thus, especially early in the game when they had not yet accrued information about the value of donuts of any given color and hole size, participants would find themselves in the position of weighting the potential positive outcomes of approaching the donut against the potential negative outcomes. The relevant outcomes include not only the potential to gain or lose points as a function of the donut's value, but also the value of gaining information about the donut by virtue of the decision to approach or missing the opportunity to obtain information by virtue of a decision to avoid. Whichever outcomes participants would weigh as more important would lead to a decision favoring either approach or avoidance. Early in the game then, participants face the task of deciding whether to approach or avoid each of the novel donut stimuli about whose value they are uncertain.

Participants were placed in one of two conditions for DonutFest. Half of the participants were given a full 5000 ms to make a decision during each trial of the game phase (unrestricted

time condition) while the other half of the participants were explicitly warned that they would have only 1000 ms to respond (restricted time condition).

Rational-Experiential Inventory. After DonutFest, participants completed the Rational-Experiential Inventory (REI; Pacini & Epstein, 1999). The REI consists of two major subscales. The first subscale, termed the Rationality scale, consists of items assessing the extent to which participants engage in a more rational way of thinking and making decisions (“I have a logical mind;” “I prefer complex problems to simple problems”). The second, termed the Experiential scale, consists of items assessing the extent to which participants engage in a more intuitive way of thinking and making decisions (“I can usually feel when a person is right or wrong, even if I can’t explain how I know;” “I often go by my instincts when deciding on a course of action”). Although the REI was originally conceived of as a trait measure, it seems possible that situational factors could increase or decrease individuals’ perception of the extent to which they utilize a certain style of thinking. It was thought that if they were relying more on a default response to make decisions in the restricted time condition, then participants in this condition would have experienced a prolonged episode of making decisions intuitively. This recent and salient experience could then lead them to indicate that they utilize intuitive thinking and decision making to a greater extent than those in the unrestricted condition when filling out the REI.

Results

REI. As hypothesized, those in the restricted condition ($M = 3.49$) indicated higher scores on the Experiential subscale ($\alpha = .89$) compared to those in the unrestricted condition ($M = 3.21$; $t(59) = 2.11, p = .04$). These results indicate that, although the REI is often conceived of

as a trait measure, restricting participants' response window increased their current self-perceived use of a more intuitive way of responding.

There were no significant differences between conditions on the Rationality subscale ($\alpha = .86$; all $ps > .1$).

Calculating the weighting bias. Following past research, the calculation of the weighting bias focused on participants' responses to the novel beans during the test phase of BeanFest (Pietri et al., 2013). In particular, the weighting bias considers the extent to which participants classify the novel beans, which resemble both positive and negative beans they had seen previously, as positive or negative *over and above* how well they learned the game-phase beans. To this end, a regression equation was used to predict participants' average response to the novel beans based on the proportion of positive and negative game beans they correctly classified during the test phase. By predicting participants' response to novel beans based on their past learning, the *residual* of this regression equation can be used as a measure of their relative weighting bias. Put in other terms, this residual measures the extent to which participants weight resemblances to a positive and resemblances to a negative *over and above* how well they learned.

Since BeanFest's inception, our laboratory has conducted multiple studies involving the same student population as participants. As a result, the data from the participants of all three experiments reported in this paper and those from a set of earlier studies were aggregated into a much larger corpus consisting of 970 participants. This collective sample allowed for a much more stable assessment of participants' relative weighting bias, less affected by the

idiosyncrasies and variability that may characterize smaller samples.³ The three variables of interest in the regression analysis were the proportion of positive game beans correct on the test phase, the proportion of negative game beans correct on the test phase, and participants' average response to the 64 novel beans. Responses to a novel bean were coded as +1 if a participant classified it as positive, and -1 if classified as negative. The average of these responses over the 64 novel beans is what is referred to as the participant's average response to novel beans. The resulting multiple regression equation from these 970 participants was as follows:

$$\text{Novel Response} = .55(\text{Positive Correct}) - .77(\text{Negative Correct}) + .07$$

Two general findings are apparent in this equation, each of which echoes observations from a multitude of past studies using the BeanFest paradigm (e.g., Fazio et al., 2004; Shook, Fazio, & Eiser, 2007; Pietri, Fazio, & Shook, 2012; Pietri et al., 2013). First, both the proportion of positive game beans correctly identified and the proportion of negative correctly identified predicted how participants responded to novel beans. This indicates that attitude generalization toward the novel beans occurred and was a function of learning: the better individuals learned the positive (negative) game beans, the more likely they were to classify novel beans as positive (negative). Second, on average, negative attitudes generalized more strongly than positive attitudes, as indicated by the larger regression weight for the negative variable relative to the positive. Indeed, using the regression equation, we see that participants who learned both bean-types equally well (e.g., 90% of both positive and negative beans correct) are expected to show

³ The reader may wonder about the extent to which the use of this collective corpus to estimate the weighting bias would correspond to regression residuals derived separately from each sample within the current article. The correlation was .94, .92, and .97 for Experiments 1-3, respectively. The strong correspondence is unsurprising given that all the studies in the collective corpus drew from the same student population. Nevertheless, the larger corpus minimizes concerns about the variability that is associated with estimates derived from any given small sample.

greater weighting of negatives over positives and, hence, classify more of the novel beans as negative (predicted response to novel beans of -.13). As this indicates, individuals showed a negativity bias on average in terms of valence weighting.

Although participants on average show a greater weighting of negatives over positives, the extent to which this occurs varies across individuals. That is, there is variability around the regression line. Responding differently to the novel beans than would be expected on the basis of how well one learned the positive and negative game beans represents deviation from the general trend revealed by the regression equation. This difference from what one would expect based on the pattern of learning, i.e., the residual, serves as the measure of an individual's valence weighting bias. Explicitly, a more positive (negative) residual indicates more weight given to positive (negative) resemblances to game phase beans than what one would expect given how well that individual learned.⁴

DonutFest behavior. Although framed as a decision of whether to approach or avoid, participants' behaviors reflected the fact that they actually had three decision choices in DonutFest: actively avoid the donut, actively approach the donut, or allow time to run out and

⁴ As the reader may guess, there may be many similar ways in which the weighting bias might be calculated using different combinations of BeanFest variables. For example, one approach would be to take a participant's average response to the game beans during the test phase and subtract it from that participant's average response to the novel beans during the test phase. A positive (negative) value would indicate responding more positively (negatively) to the novel beans than the game beans and therefore showing a positive (negative) weighting bias. This approach results in a correlation with the current residual method of .94 across the present three experiments and shows nearly identical results with those reported in all three experiments. The current residual approach is preferred as it both builds on previously-conducted research and better highlights the phenomenon of attitude generalization in a direction, and to an extent, that is greater than expected on the basis of the pattern of learning exhibited by the participant.

thereby avoid the donut (as a lack of a response resulted in no gain of information or point change). As is to be expected, participants in the restricted condition ($M = 23.81$, $SD = 15.21$), who had only 1000 ms to respond, tended to miss the response deadline on many more of the 120 trials than those in the unrestricted condition ($M = .80$, $SD = 1.21$, $t(59) = 8.40$, $p < .001$). For this reason, we focused first on those trials on which participants expressed an active decision. We calculated the proportion of approach decisions each participant made relative to the number of donuts on which the participant expressed a decision in the allotted time.

Our major hypothesis is that decisions to approach versus avoid will reflect individuals' valence weighting tendencies to a greater extent when the time to deliberate is restricted. However, as noted earlier, we also anticipated that individuals' weighting biases may be most evident early in the game prior to their learning the value of the different types of donuts, i.e., when the environment was still novel. In particular, we reasoned that in later trials of DonutFest, participants should have less need to construct behavioral decisions on the basis of valence weighting tendencies since they have likely learned the actual valence of the donuts. Hence, our primary prediction is for a two-way interaction between weighting bias and condition, with the possibility that this effect may be moderated by the DonutFest block variable.

As block is a within-subjects measure, we used multilevel modeling to test our hypotheses. What follows are the details of how the model was constructed. In this case, the three blocks of DonutFest were nested within participants at level 1 while the weighting bias and condition were at level 2. Condition was effects coded, blocks were coded linearly, and weighting bias was grand-mean centered. Robust standard errors were assumed and based on a recommendation by Nezlek (2011), effects were only estimated as random if they were significant at the 0.2 level or less.

As predicted, there was a significant two-way interaction between weighting bias and condition ($\gamma = 0.27$, $t(61) = 2.43$, $p = .02$; see Figure 1) indicating that, on average, those with a positive weighting bias actively approached a greater proportion of donuts than those with a negative weighting bias in the restricted condition ($\gamma = 0.40$, $t(61) = 2.33$, $p = .02$), but not the unrestricted condition ($\gamma = -0.14$, $t(61) = 0.33$, $p = .33$). There was also tentative support for our secondary hypothesis that the difference in behavior should be most apparent in earlier rather than later trials. This was evidenced by a marginal three-way interaction between weighting bias, block, and condition ($\gamma = -0.07$, $t(61) = 1.66$, $p = .10$). As predicted, this effect indicated that the weighting bias by condition effect was strongest in Block 1 of DonutFest ($\gamma = 0.23$, $t(61) = 2.48$, $p = .02$), but then weaker in Block 2 ($\gamma = 0.06$, $t(61) = 0.57$, $p = .57$) and Block 3 ($\gamma = 0.07$, $t(61) = 0.99$, $p = .33$).⁵ Simple slopes analyses of Block 1 confirmed that those with a more positive weighting bias were approaching a greater proportion of donuts than those with a more negative weighting bias in the restricted condition ($\gamma = 0.34$, $t(61) = 2.37$, $p = .02$), but not the unrestricted condition ($\gamma = -0.11$, $t(61) = 1.00$, $p = .32$).

⁵ As we would assume, participants did learn across the three DonutFest blocks. For each block, we computed the proportion of donuts correctly classified for each participant. If participants are learning we should see this proportion becoming larger for each progressive block. A 2 (condition) X 3 (block) mixed ANOVA revealed two main effects. First, participants in the unrestricted condition ($M = .75$) learned better than those in the restricted condition ($M = .54$; $F(1, 59) = 39.61$, $p < .001$). This finding makes sense as the unrestricted participants were not rushed and therefore had more time to think and strategize about their decisions. Second, and just as anticipated, learning increased across blocks with Block 1 showing the smallest proportion correct ($M = .60$), followed by Block 2 ($M = .65$), and then Block 3 ($M = .68$; $F(2, 118) = 14.91$, $p < .001$). There was no block by condition interaction ($F(2, 118) = .05$, $p = .82$).

We can also break this analysis down further to examine whether the weighting bias seems to be most related to avoids, approaches, or to both types of decisions. Using the same multilevel model as described above, we first examined the total number of active avoid decisions. As with the proportion of donuts approached, we found the predicted weighting bias by condition interaction when predicting decisions to avoid ($\gamma = -12.27$, $t(61) = 3.04$, $p = .003$; see Figure 2). This interaction indicated that those with a more negative weighting bias tended to make more active decisions to avoid than those with a more positive weighting bias in the restricted condition ($\gamma = -19.18$, $t(61) = 3.28$, $p = .002$), but not in the unrestricted condition ($\gamma = 5.36$, $t(61) = 0.96$, $p = .34$). The three-way interaction between weighting bias, block, and condition also was significant ($\gamma = 3.19$, $t(61) = 1.96$, $p = .05$). As before, this result indicated that the weighting bias by condition interaction was strongest in Block 1 ($\gamma = -10.42$, $t(61) = 3.39$, $p = .003$), and weaker in Block 2 ($\gamma = -3.20$, $t(61) = 3.66$, $p = .39$) and Block 3 ($\gamma = -4.05$, $t(61) = 1.08$, $p = .31$). Simple slopes analyses regarding the Block 1 interaction confirmed that those with a more negative weighting bias avoided donuts to a greater extent than those with a more positive weighting bias in the restricted condition ($\gamma = -16.17$, $t(61) = 3.22$, $p = .002$), but not the unrestricted condition ($\gamma = 4.67$, $t(61) = 1.03$, $p = .31$).

A similar weighting by condition interaction was not observed with respect to the number of approaches to novel donuts ($\gamma = 3.63$, $t(61) = 0.78$, $p = .44$). There was also no weighting bias, block, and condition interaction ($\gamma = -2.14$, $t(61) = 1.59$, $p = .18$). It appears, then, that the interactions for the proportion of donuts approached were driven mostly by participants' decisions to actively avoid the novel donuts or not.

Recall that participants, especially those in the restricted condition, sometimes failed to make an active decision within the allotted time. It seems possible the difficult requirement of

coming to a decision in under 1000 ms may have led some participants to assume a pattern of allowing time to run out if they were not sufficiently interested in approaching a given donut.

To investigate this possibility, we predicted the number of missed trials following the method above. As with those analyses above, there was a significant weighting bias by condition interaction ($\gamma = 8.65$, $t(61) = 2.37$, $p = .02$). This two-way interaction revealed that those with a positive weighting bias missed more trials than those with a negative weighting bias in the restricted condition ($\gamma = 15.90$, $t(61) = 2.22$, $p = .03$), but not in the unrestricted condition ($\gamma = -1.40$, $t(61) = 1.08$, $p = .28$). There was no significant weighting bias, block, and condition interaction simply indicating that this relationship existed across each block in the restricted condition and may indeed represent a pattern of behavior participants had assumed in light of the strict time limit.

The question may be asked as to why those with a more positive weighting bias tended to miss more trials than those with a negative weighting bias. It is possible that participants with a more positive weighting bias may have done so as a by-product of their focus on approaching donuts. It may be that while it was of primary importance for those with a negative weighting bias to be vigilant for those novel stimuli they wished to avoid – as reflected in their increased avoidance behavior – those with a positive weighting bias may have concentrated on approaching novel stimuli. In concentrating on approaching rather than avoiding, participants with a positive weighting bias may have simply allowed time to run out when they did not wish to approach. Indeed, when participants did not respond on a trial, it was functionally equivalent to avoiding; in both cases, no further information was acquired and points were not gained or lost. Thus, while those with a negative weighting bias may have been concerned with actively

avoiding, those with a positive weighting bias may have attended more singularly to approach-oriented decisions.

Latencies to actively approach versus avoid. If it is the case that those with a more positive weighting bias are focusing relatively more on approaching while those with a negative weighting bias are focusing more on avoiding, we reasoned that the latencies with which participants in the restricted time condition actively approached versus actively avoided donuts may provide additional evidence for this focus. Specifically, those with a more positive weighting bias may show faster response times when actively approaching a donut than when actively avoiding a donut. To this end, a difference score was calculated for each block in the restricted time condition as the average response time, in milliseconds, when a participant approached a donut subtracted from the average response time it took to avoid a donut. Averaged across all blocks, the mean score on this index was significantly less than zero, ($M = -47.52$, $SD = 73.07$, $t(30) = 3.62$, $p = .001$), indicating that participants were generally faster to avoid than approach.

We predicted this difference score from the weighting bias, block, and their interaction.⁶ As predicted, there was an overall effect of the weighting bias ($\gamma = 173.78$, $t(24) = 2.34$, $p = .03$). Thus, the general tendency to make approach decisions more quickly than avoidance decisions was more characteristic of those with a more positive weighting bias.).⁷ There was no block by

⁶ All participants were represented in Block 1, but scores were lacking for four participants in Block 2 and five in Block 3 because they failed to make any active decisions to avoid. Instead, they permitted time to run out on various trials, apparently as a means of avoiding. The score of an additional participant was excluded in Block 2 because it was over three standard deviations from the mean.

⁷ Including the unrestricted condition in these analyses does not appear appropriate in the present case, because there was a marked difference in standard deviations for the latencies between the unrestricted ($SD = 151$ ms) and the

weighting bias interaction ($\gamma = 5.13$, $t(24) = 0.11$, $p = .92$) indicating that the relationship existed across blocks.

Discussion

The results from this first experiment show that the weighting of positives and negatives has a larger impact under time pressure and this, in turn, provides general support for the idea that behavioral decisions begin with the construction of an initial evaluative response that is influenced by individuals' valence weighting tendencies. Those participants who had little time to deliberate beyond their initial default response indicated more intuitive decision making and thinking on the REI items. More importantly, when they had little time, participants displayed greater reliance on their weighting bias while navigating the novel DonutFest environment. When looking at trials on which participants made an active decision, those with a more positive weighting bias approached a greater proportion of donuts than those with a more negative weighting bias. Breaking this proportion score into its constituent parts revealed that this relationship was driven by avoidances: those with a more negative weighting bias actively avoided more novel stimuli compared to those with a more positive weighting bias. The results for the proportion of donuts approached and the total number of donuts avoided also supported the hypothesis that the weighting bias should have its largest effect in novel situations as earlier trials in DonutFest tended to show the strongest relationship between behavior and the weighting bias. Indeed, as they learned of each type of donut, participants could base their subsequent decisions on their actual knowledge rather than valence weighting.

restricted ($SD = 73$ ms) conditions. This renders comparisons of response times across the two conditions rather dubious. For the interested reader, there were no significant effects of weighting bias, block, or their interaction on response time in the unrestricted condition ($ps > .60$).

There was also a third option in DonutFest whereby individuals could effectively avoid a donut by allowing time to run out. Participants did not fail to respond at random. In the restricted time condition, those with a more positive weighting bias missed more trials. A more singular focus on approaching may have led such individuals to allow time to run out on a given trial, thereby effectively avoiding the donut. This argument is supported by the time it took rushed individuals to make decisions in line with their weighting bias. They were relatively quicker to respond when making bias-congruent (e.g., positive weighting bias making quicker approach decisions) as opposed to bias-incongruent decisions (e.g., positive weighting and avoiding). Thus, when making the decision of whether to approach or avoid the novel stimuli, those with a negative weighting bias gave greater weight to the potential negative outcome of losing points while those with a more positive weighting bias gave greater weight to the positive outcome of learning about the environment and the possibility of earning points.

Obviously, the weighting bias was not as pivotal in the unrestricted time condition as in the restricted condition. This outcome likely stemmed from the unrestricted participants' strategic deliberations regarding their approach and avoidance decisions. When given extra time, participants had the opportunity to deliberate and strategize regarding their decisions and, hence, exhibited better learning of donuts' valences (see Footnote 5). They could, for example, spend more time considering the extent to which a currently presented donut resembled one they had seen previously. The weighting bias likely exerted some initial influence on these participants' initial appraisal of whether to approach or avoid a given donut, but those appraisals would be updated and modified as participants in the unrestricted condition utilized their developing theories and predictions regarding the visual appearances of positive versus negative donuts.

Overall, the time restriction utilized in this experiment enhanced the strength of the relation observed between individuals' weighting bias scores and their sampling behavior in a novel environment; this was particularly true on earlier trials when participants had acquired little direct knowledge on which to rely. On the basis of these results, the weighting bias seems to be most robust under time pressure and can be considered to play a role in the construction of an initial default response.

Experiment 2

After restricting participants' opportunity to override their weighting bias response, we also wanted to manipulate a second variable postulated to be important by the MODE model, namely participants' *motivation* to follow or deliberate further about the initial evaluative response constructed as a function of their weighting bias. Based on the tenets of the MODE model, we hypothesized that giving participants reason to override their weighting bias should lead to less correspondence between the weighting bias and behavior while giving participants reason to follow their weighting bias should lead to greater correspondence.

In this experiment, we utilized the Balloon Analogue Risk Task (BART; Lejuez et al., 2002), which is considered a proximate measure of risk-taking behavior. The BART measures how many times participants are willing to pump up a virtual balloon in order to increase its value. Participants have to weigh the risk of popping the balloon and receiving no points from that round, against pumping the balloon fewer times, but also earning fewer points. Past research has indicated that self-reported risk-taking in health and safety domains is related to greater risk-taking in the BART (Lejuez et al., 2002; Lejuez, Simmons, Aklin, Daughters, & Dvir, 2004; Lejuez et al., 2003). Further research has shown that the more positive an individuals' valence weighting bias, as assessed via the BeanFest procedure, the more individuals pump the balloons

(Pietri et al., 2013). These findings indicate that those with a positive weighting bias may be giving greater weight to the relative positive outcome of earning incrementally more points compared to the relative negative outcome of popping the balloon and losing points.

Explicitly stated, then, when participants are motivated to follow their bias, we expected to replicate the original finding that the more positive individuals' weighting bias, the more they should be willing to pump up the balloon and risk it popping. When participants were motivated to override their bias, however, this relationship should be attenuated. It is precisely this differential motivation that we manipulated in Experiment 2.

Method

Overview. Participants entered the lab and were told they were participating in three separate studies. For the first study, participants were told they would be playing a game called BeanFest in order to see how individuals navigate novel environments. As in Experiment 1, we used BeanFest to measure participants' weighting bias. After receiving an incomplete debriefing about BeanFest, participants were re-consented for what they believed was the second study. During this portion of the experiment, participants read three newspaper articles and were asked to rate them on different aspects. This part of the experiment constituted the manipulation. For half of the participants, the final article they read was a summary of scientific research reporting the findings that *following one's intuition* leads to better life outcomes. The other half of the participants read a similar final article, but one that reported that *overriding or overcoming* intuition leads to better life outcomes. After again receiving an incomplete debriefing and then being re-consented, participants completed the BART. Finally, participants completed the REI as in Experiment 1.

Participants. Participants were 61 (35 male and 26 female) introductory psychology students who completed the experiment for class credit. Two participants were excluded for pumping fewer than three standard deviations below the mean on the BART. A total of 59 participants remained for subsequent analyses (33 male and 26 female).

Procedure.

BeanFest: Assessing individuals' weighting bias. As with Experiment 1, participants played BeanFest in order to obtain behavioral estimates of their weighting bias. The equation from the aggregated sample of 970 participants was again used to calculate individuals' relative valence weighting bias.

Motivation manipulation. Participants were told that this second portion of the experiment was in collaboration with a newspaper aiming to assess the kinds of psychology articles students their age find most appealing. To this end, participants read three different articles and rated each on different characteristics (e.g., how entertaining they were). In fact, the first two articles – an article on autism and an article on sleep deprivation – were held constant across the two conditions while the final article differed based on condition.

For half of the participants, the final article, which was attributed to *The New York Times*, was a 255-word summary of ostensible scientific research showing that following one's intuition or gut-instinct leads to a longer, healthier, and more successful life ("follow" condition). The other half of the participants read a parallel article summarizing the research as showing that overriding one's intuition or gut-instinct leads to these same benefits ("override" condition). In the former case, the article's headline read: "Trusting gut-reactions leads to the best decisions," and in the latter: "Overcoming gut-reactions leads to the best decisions."

BART: The dependent variable. In the BART, participants were presented with 20 trials where they were required to decide how many times they wished to pump up a virtual balloon. Each time they wanted to pump the balloon, participants clicked a button on the lower left-hand of their screen labeled “Pump up the balloon.” Each time they pumped up the balloon, participants potentially earned virtual money (5¢).⁸ This money was not added to their total, however, if they did not press the “Collect” button before the balloon popped. Thus, participants needed to constantly weight the relative negative outcome of popping the balloon and losing all of their money on that trial against the relative positive outcome of pumping the balloon one more time to increase its value. After popping the balloon or collecting their earning for that trial, a new balloon was displayed and participants played the game again with this new balloon. The maximum number of pumps was set to 25 for each balloon. The a priori probability of a balloon exploding on any given pump was 1/25. Thus, unlike DonutFest in Experiment 1, the BART does not allow much in the way of learning over the course of the task as balloon pops occur largely at random. Due to this, we did not anticipate the effect of weighting bias to differ across the trials.

REI. The REI questionnaire provided an effective manipulation check in Experiment 1. Hence, we employed it again in the present experiment. Participants completed the REI upon finishing the BART. The Rationality ($\alpha = .89$) and Experiential ($\alpha = .92$) subscales again showed good reliability.

Results

⁸ Participants did not actually earn any real money and were told in the instructions that they would just be playing for virtual money.

REI. If they were relying less on a default response to make decisions in the “override” condition, then participants in this condition may have indicated that they are able to and actually utilize intuitive thinking and decision making *less* than those in the “follow” condition. Those in the “override” condition ($M = 3.07$) did indeed indicate lower scores on the Experiential subscale compared to those in the “follow” condition ($M = 3.66$, $t(57) = 3.73$, $p < .001$).

There was also some indication that the manipulation affected the extent to which participants reported using rational thinking in decision making, with those in the “override” condition ($M = 3.85$) tending to report greater rational thinking than those in the “follow” condition ($M = 3.60$, $t(57) = 1.61$, $p = .11$).

Thus, the manipulation appears to have been successful in decreasing participants’ self-reports of using a more intuitive thinking style and, to some extent, increasing their use of a more rational thinking style.

BART. Do these changes carry over into actual behavior? The answer appears to be “yes.” Following Lejuez and colleagues’ (2002) recommendation, we focused on pumping behavior as the dependent measure. As anticipated, participants’ weighting bias was related to pumping behavior, but only in the “follow” condition. Specifically, the more positive their weighting bias, the more times participants pumped the balloon in the “follow” condition (see Figure 3). These results were obtained using a multiple regression equation predicting average pumps per trial from condition (“follow” vs. “override”), weighting bias, and their interaction. All variables reported were standardized and the conditions were effects coded. As predicted, there was an effect of weighting bias ($B = .45$, $t(55) = 1.96$, $p = .06$) indicating that, regardless of condition, the more positive participants’ weighting bias, the more they pumped. The only other significant effect was the predicted condition by weighting bias interaction ($B = -.56$, $t(55) =$

2.42, $p = .02$). A simple slopes analysis revealed that those with a more positive weighting bias pumped more than those with a more negative weighting bias when in the “follow” condition ($B = 1.01$, $t(55) = 2.82$, $p = .01$), but not the “override” condition ($B = -.11$, $t(55) = .36$, $p = .72$).⁹

Discussion

The results from this second experiment provide additional support for our hypothesis that individual differences in the weighting of positives versus negatives play a role in the construction of an initial default response that can be used in decision making. Those participants who were motivated to override their initial responses did not show a relationship between the weighting bias and balloon pumps, whereas those who were motivated to follow their default showed the hypothesized relationship. Given that past research has observed correspondence between the valence weighting measure and pumping behavior (Pietri et al., 2013), we are not in a position to draw any inference as to whether motivating participants to rely on intuitive thinking actually increased the expected correspondence. However, it is readily apparent that participants did override their initial bias when motivated to do so.

As with the first block of the DonutFest paradigm, the BART also provides very little information and represents an environment where participants can have few expectations or theories to inform their decision making. Indeed, as in Experiment 1, it appears that participants do not make decisions at random when given little information, but instead may rely on the extent to which they weight positives over negatives or vice versa. As discussed earlier, it is precisely in these situations where the weighting bias should have its largest impact on decision making.

⁹ As anticipated, there was no three-way interaction between trial, weighting bias, and condition ($p = .34$), indicating that the effect of weighting bias did not change as the trials progressed.

Experiment 3

Experiments 1 and 2 have demonstrated the role of the valence weighting bias in developing an initial default response. To illuminate its importance further, we can also ask about the consequences of acting on the basis of responses influenced by one's weighting bias. As shown in Experiments 1 and 2, individuals with a more negative weighting bias tended to give more weight to the potential to lose points, avoided engaging with novel stimuli, and in that sense tended to explore their environment less fully. On the other hand, those with a more positive weighting bias tended to give more weight to potential rewards, to therefore engage relatively more with the novel environment, and, particularly in Experiment 2, tended to test their environment to a greater extent through greater pumping.

One consequence of these valence weighting proclivities could be that initial positive or negative information (e.g., culturally-transmitted prejudices) may overly affect individuals with a given weighting bias due not only to their immediate effect on attitude formation, but also as a function of their downstream impact on subsequent behavior. For instance, if individuals with a negative weighting bias were to initially encounter information that was negative, they may be all the more likely to heed this information and subsequently avoid whatever stimuli were associated with that negativity. If the negative information was actually incorrect, failing to explore these stimuli fully could lead to the maintenance of the mistaken view (Fazio et al., 2004). Conversely, those with a positive weighting bias may give less weight to this negative information, engage with the stimuli to a greater extent, and thereby gain a potentially more accurate view.

Predictions for initially encountering false positive information are less clear. On the one hand, those with a more positive weighting bias may give this positive information more

credence, paralleling our argument that those with a negative weighting bias should heed negative information more strongly. However, another possibility is that, regardless of their weighting bias, individuals may heed the positive information sufficiently to engage with the stimuli. Through such exploratory behavior, all individuals, including those with a negative weighting bias, will obtain experiential feedback about the true value of the stimuli and therefore come to similar understandings of them. This latter prediction is consistent with the evidence obtained by Fazio et al. (2004) of a striking difference between invalid positive versus invalid negative prejudices. Because the former encourage approach behavior, whereas negative prejudices promote avoidance, individuals can overcome initial invalid positive beliefs, but are likely to maintain invalid negative beliefs. Whereas Fazio et al. demonstrated this general pattern, our current argument is that it will prove especially true of those individuals characterized by a negative weighting bias.

To test these hypotheses, an environment was constructed in which all stimuli were equally positive. Before exploring these stimuli, however, individuals were given false valenced information about different types of stimuli in an attempt to prejudice them as a function of their weighting bias. Thus, unlike Experiment 1, in which individuals' decisions to approach or avoid novel stimuli were not biased by initial prejudicial information, Experiment 3 sought to manipulate initial expectancies.

Method

Overview. As in Experiments 1 and 2, participants' weighting biases were estimated through the BeanFest paradigm. After completing BeanFest, participants played a variation of the BART from Experiment 2. In the current version, the balloons varied in color and participants were provided with invalid positive information regarding the strength of one color

balloon and invalid negative information regarding another color. This information was meant to prejudice participants' initial attitudes toward the two balloon types, despite the fact that they did not actually differ in strength. After playing the BART game, participants provided assessments of each balloon type.

Participants. Participants were 66 (24 male and 42 female) introductory psychology students who completed the experiment for class credit.

Procedure.

BeanFest: Assessing individuals' weighting bias. As with Experiments 1 and 2, participants played BeanFest in order to obtain behavioral estimates of their weighting bias. The equation from the aggregated sample of 970 participants was again used to assess individuals' weighting bias.

BART. The BART procedure in the current experiment was similar to that in Experiment 2, but with a few key differences. In this variation on the game, there were now five different balloon colors – aqua, blue, orange, purple, and yellow. Furthermore, prior to playing the actual BART game, participants were given information pertaining to two of the balloon colors. As a cover story, participants were told that the research concerned the transmission of cultural knowledge. To this end, participants were asked to randomly select a folder that contained information from a past generation of players about the upcoming BART game they would be playing. Although participants believed they were selecting information from a past generation at random, all of the folders contained the same information.

In each folder were handwritten notes from two past players indicating that each player had done very well on the task. In response to a question on the form that invited suggestions regarding the balloon game, one of the players advised that one particular balloon color seemed

very strong and, hence, a “good” balloon, whereas the other player stated another color seemed very weak and, hence, a “bad” balloon. This information thereby attempted to create initial attitudes toward two of the balloon colors that were either relatively positive or negative. No information was given about the remaining balloon colors. These “neutral” balloons served as a control, permitting us to compare the consequences of receiving false positive or false negative information to a relevant baseline.

For the sake of counterbalancing, one half of the participants received information that the aqua balloons were the “good” balloons and the purple were “bad,” and vice versa for the other half of the participants. Coding for the counterbalanced conditions did not yield differing results and therefore will not be discussed in subsequent analyses. The experiment thus involved a within-subjects design with all participants receiving information about two different balloon colors.

Importantly, although past generations indicated that there were “good” and “bad” balloons, all balloons actually had the same exact probability of popping. As in Experiment 2, the a priori probability of any balloon popping was $1/25$ on a given trial. There was a total of 41 balloon trials, with the “good” and “bad” balloons each being shown 10 times and each of the three filler balloon colors shown 7 times. The order in which the balloons were displayed was arranged in one of two pseudo-random sequences. Each sequence involved a randomized order that was then modified so that multiple balloons of the same type would not cluster together. Coding for the two sequences did not reveal differing results. All other BART parameters were the same as in Experiment 2.

Balloon assessments. After the BART, participants completed a questionnaire regarding their assessments of each balloon type. Participants were asked on a 0 to 8 scale (not at all;

extremely) how dangerous, threatening, safe, unsafe, hazardous, perilous, unpredictable, erratic, positive, and negative balloons of each color were. Negative items were reverse-coded so that higher numbers indicated more positive attitudes. These judgments were then averaged to create general attitudes toward the different balloon types. Scales for each balloon type showed good reliability ($\alpha > .84$).

Self-reported weighting bias. In the interest of further examining the earlier-noted finding regarding a null relation between the performance-based measure of valence weighting and direct self-reports, participants also completed the Weighting Bias Questionnaire (WBQ; Pietri et al., 2013). The WBQ consists of four very direct questions asking participants to self-report the extent to which they believe they weight positive versus negative information (e.g., “To what extent do you tend to give more weight to positive information over negative information?”) on a 7-point scale. More positive numbers indicate more weight given to positives ($M = 4.29$, $SD = .89$, $\alpha = .58$).¹⁰

Results

Manipulation check. The first question we can ask is simply whether or not the information participants received about the balloons was effective in biasing their behavior in the BART game. In order to index behavior during the BART, the average number of pumps per trial a participant gave to the “good,” “neutral,” and “bad” balloon colors was calculated separately. A repeated-measures ANOVA revealed a significant effect of balloon type on pumping behavior ($F(2, 130) = 37.47$, $p < .001$). Planned comparisons revealed that the means

¹⁰ The current alpha of .58 is inexplicably lower than the .73 obtained in Pietri et al. (2013). Deletion of one item within the WBQ led to a higher alpha of .66 for the current study. Analyses using this composite yielded outcomes equivalent to those reported in the text.

for all three balloon types differed from one another, all $ps < .001$, with the “good” being pumped the most ($M = 9.12$), followed by “neutral” ($M = 7.27$), and finally the “bad” ($M = 5.95$).

Weighting bias and pumping behavior. Though our manipulation was successful on average, our hypothesis was that some types of valenced information would be more or less important for individuals with a given weighting bias. Did, in fact, the impact of the initial information about the “good” versus “bad” balloons vary as a function of individuals’ weighting biases?

To test this hypothesis, a difference score was created subtracting the average number of pumps to the “bad” balloons from the average number of pumps to the “good” balloons. Higher numbers indicate more pumping on the “good” balloon relative to the “bad.” Using linear regression we then employed individuals’ weighting bias to predict differential pumping behavior as a function of the initial information provided about the two balloons. All variables reported were unstandardized. Confirming our hypothesis, we found that a more negative weighting bias was associated with greater differentiation ($b = -5.71$, $t(64) = 2.39$, $p = .02$). In other words, whereas a more negative weighting bias was associated with greater differentiation between the balloon types, those with a more positive weighting bias pumped the balloons more equally. Indeed, at one standard deviation below the mean on weighting bias – a more negative weighting bias – individuals pumped on average 4.20 times per trial more on the “good” than the “bad” balloons. Individuals with a more positive weighting bias – one standard deviation above the mean – however, showed nearly half as much differentiation ($M = 2.14$). Both difference scores differed from zero – the point of equivalent pumping, $ps < .001$ – indicating that both those with more positive and more negative weighting biases heeded the initial information, but to a different extent.

To break this analysis down further, two additional difference scores were examined to compare the effects of the initial valenced information to stimuli about which participants received no information. This was done by subtracting the average pumps given to the “neutral” balloons from the average pumps given to the “good” balloons and subtracting the average pumps to the “neutral” from the “bad.”

Consistent with our earlier speculations about the impact of initial positive versus negative information in the BART, the “good” versus “bad” difference score results appeared to be primarily attributable to differential pumping on the “bad” balloons. Weighting bias scores were unrelated to the difference score comparing the “good” balloon type to the “neutral” ($b = -.81$, $t(64) = .44$, $p = .66$), but were associated with differential pumping to the “bad” compared to the “neutral” ($b = 4.90$, $t(64) = 2.86$, $p < .01$; see Figure 4).¹¹ Specifically, despite the previous generation’s admonitions, those with a more positive weighting bias (1SD above the mean) seemed not to heed the negative advice about the “bad” balloons and pumped them similarly to how they pumped the “neutral” balloons, about which they had received no information ($M = .44$). A one-sample t-test comparing this mean difference to zero confirmed that, indeed, those with this more positive weighting bias did not pump differentially between the “bad” and “neutral” balloon types ($t(64) = 1.00$, $p = .32$). In contrast, those with a more negative weighting bias (1SD below the mean) took the warning seriously and pumped less on the “bad” compared

¹¹ For the purpose of illustrating each balloon type’s relation to the weighting bias graphically, three regression equations were used to predict pumping behavior from individuals’ weighting biases. The following linear relations emerged and are graphed in Figure 5: “good,” $b = 1.28$, $t(64) = .68$, $p = .50$; “neutral,” $b = 2.09$, $t(64) = 1.32$, $p = .19$; “bad,” $b = 6.99$, $t(64) = 3.66$, $p < .001$.

to the “neutral” ($M = -2.20$). This mean difference score was significantly different from zero ($t(64) = 5.07, p < .001$).¹²

Weighting bias and balloon assessments. The next question we can ask is whether the influence of the initial information given to participants was also evident on their post-task assessments of the balloons and whether any such influence was associated with the weighting bias. A repeated-measures ANOVA revealed that, on average, at the end of the BART, participants rated the balloons differently from one another ($F(2, 130) = 10.37, p < .001$). Planned comparisons revealed that the means for all three balloon types differed from one another, all $ps < .04$, with the “good” being rated most positive ($M = 6.02$), followed by “neutral” ($M = 5.46$), and then the “bad” ($M = 4.89$).

To test whether the weighting bias was also related to post-task evaluations of the balloons, a difference score was again created by subtracting the average evaluation of the “bad” balloons from that for the “good.” Using linear regression to predict this difference score from the weighting bias, it was found that a more positive weighting bias was associated with showing less differentiation in attitudes between “good” and “bad” balloons ($b = -3.44, t(64) = 2.02, p = .05$). In other words, a more positive weighting bias (1SD above the mean) was associated with

¹² As indicated by Footnote 9, we did not find a significant effect of weighting bias on the “neutral” balloons, though the direction of the effect is in the hypothesized direction. Those with a more positive weighting bias tended to pump more on the “neutral” balloons than those with a more negative weighting bias $b = 2.09, t(64) = 1.32, p = .19$. In our view, that there is no statistically significant relation for the “neutral” balloons should not be taken as inconsistent with the results from Experiment 2. Unlike Experiment 2, Experiment 3 used multiple balloon colors, which participants were led to believe differed in strength. This would allow participants to take a more nuanced approach when determining pumping behavior; indeed, they were essentially encouraged to actively strategize about the different balloons types.

evaluating the balloons more equivalently by the end of the task ($M = .51$), whereas a more negative weighting bias (1SD below the mean) was associated with more disparate evaluations ($M = 1.75$). Indeed, when tested against zero – representing no differentiation between the balloon types – those with a more positive weighting bias showed no significant differentiation between the balloon types ($t(64) = 1.19, p = .24$) while those with a negative weighting bias did ($t(64) = 4.06, p < .001$). Given that there was actually no difference between the balloon types, this means that those with a more positive weighting bias were uncovering the truth, whereas those with a more negative weighting bias continued to believe the invalid information that they received at the beginning of the task.

As with pumping behavior, two additional difference scores were created comparing attitudes toward the “good” and “bad” balloons to the “neutral.” Weighting bias was associated with differential attitudes toward the “bad” compared to “neutral” balloons ($b = 2.31, t(64) = 2.03, p = .05$; see Figure 5).¹³ Those with a more negative weighting bias (1SD below the mean) maintained the invalid negative view that had been transmitted by their earlier-generation partner ($M = -.99; t(64) = 3.43, p < .001$), continuing to believe that the “bad” balloons were worse than the balloons about which they had received no information. In contrast, those with a more positive weighting bias overcame this initial invalid information and rated the balloons equivalently ($M = -.16; t(64) = .55, p = .59$). No such relation was observed for the difference score involving attitudes toward the “good” balloons relative to the “neutral” balloons ($b = -1.13,$

¹³ Again, three regression equations were used to illustrate individuals’ final attitudes toward each balloon type as a function of their weighting biases: “good,” $b = -1.30, t(64) = 1.11, p = .27$; “neutral,” $b = -.17, t(64) = .36, p = .72$; “bad,” $b = 2.15, t(64) = 2.06, p = .04$.

$t(64) = .95, p = .35$). Weighting bias scores were unrelated to the extent to which participants differentiated between “good” versus “neutral” balloons.

Mediation analyses. Presumably, individuals’ post-task attitudes toward the balloons were affected by their experience with the balloons during the game. That is, the weighting bias was thought to have its effect on post-task attitudes through its more proximal effect on in-game pumping behavior. Specifically, those with a more positive weighting bias were hypothesized to give less weight to the initial negative information, to therefore pump more similarly on each balloon type, and to thereby show less post-task differentiation between the two balloons. Those with a more negative weighting bias were hypothesized to give more weight to the initial negative information and to then deliver relatively few pumps to the “bad” balloon. As a result, they would not learn the balloon’s actual strength, not discern that it was stronger than they had been led to believe, and to thereby maintain their invalid negative attitudes toward this balloon type.

To examine whether differential pumping behavior mediated the relationship between the weighting bias and differential final attitudes, a mediation analysis using bootstrapping with 20,000 observations was utilized (Preacher & Hayes, 2004). Results were consistent with the proposed mediation model as the bootstrapping 95% confidence interval did not include zero $[-4.73, -.83]$ (see Figure 6). In other words, the data are consistent with a causal model suggesting that a more positive weighting bias led individuals to disregard the initial negative information provided to them, to then pump more similarly on both the “good” and “bad” balloons, which then led them to show less of a difference in final attitudes toward the balloons. Conversely, the mediational model implies that a more negative weighting bias led individuals to heed the initial negative information provided to them, to then pump less on the “bad” balloons

compared to the “good,” which then led them to maintain their initial attitudes toward the balloons.

When analyzing the “bad” versus “neutral” balloons separately using the bootstrapping procedure, the same pattern of results was obtained such that differential pumping between these two balloon types mediated the relationship between the weighting bias and final differential attitudes [.35, 3.60].¹⁴ For the “bad” compared to “neutral” balloons, the weighting bias affected final attitudes through pumping behavior.

These analyses indicate that everyone, regardless of weighting bias, heeded the initial positive information, but only those with a positive weighting bias pumped the “bad” balloons sufficiently to come to the realization that the “bad” and “neutral” balloons did not differ. Those with a negative bias were mindful enough of the negative advice that they never tested the bad balloons sufficiently and, hence, continued to view them as different from the “neutral” despite them being equivalent in nature.

Weighting Bias Questionnaire. As has been reported in previous BeanFest work, participants appear unable to report their valence weighting tendencies as there was no correlation between the weighting bias measured in BeanFest and the self-report WBQ ($r(64) = .08, p = .52$; see Pietri et al., 2013).

The WBQ also did not predict any of the various balloon variables by itself or in interaction with the weighting bias ($ps > .10$). When both the WBQ and the weighting bias were entered simultaneously into a regression equation, the weighting bias continued to be significantly predictive of the balloon variables ($ps < .05$).

¹⁴ The “good” compared to “neutral” difference score was not tested for mediation as there was no initial effect of weighting bias on pumping behavior for this score.

General Discussion

The central aim in the present research was to better illuminate the weighting bias and when it may be most impactful. Across the first two experiments, substantial evidence was obtained that the extent to which individuals show a proclivity to weight the positive versus negative aspects of stimuli may be most impactful in novel situations characterized by time pressure or a lack of motivation to deliberate further. Both of the dependent measures employed in the experiments, approaching or avoiding novel donuts in DonutFest and pumping balloons in the BART, emanate from novel environments where participants had no previous knowledge. Given this lack of knowledge, participants made decisions that showed evidence of their valence weighting tendencies when they either had little time to do otherwise, or when they were not motivated to deliberate. In other words, individuals appear to construct an initial response to novel or ambiguous situations that reflects their valence weighting bias, and they are content to treat this response as a default basis for action, unless the situation provides them with both the motivation and opportunity to do otherwise.

In Experiment 1, participants with a more positive weighting bias who had little time to override their default response approached a greater proportion of novel stimuli when they actively made a decision. Furthermore, when participants had a more negative weighting bias, they avoided novel stimuli to a greater extent. Finally, it was also found that individuals were quicker to respond when making bias-congruent decisions (e.g., positive weighting bias making quicker approach decisions) as opposed to bias-incongruent decisions (e.g., positive weighting bias and avoiding). In Experiment 2, participants who were motivated to override their default response showed no correspondence between their weighting bias and our behavioral risk-taking measure. Those who were motivated to follow their default response showed a strong

relationship between their weighting bias and their propensity for risk-taking. Specifically, those with a more positive weighting bias took more risks than those with a negative weighting bias. Both experiments' results are in line with our reasoning from the perspective of the MODE model, indicating that when they have both sufficient motivation and opportunity, individuals will tend to use a more deliberative processing style and rely less on any default response they may have activated or constructed (Fazio, 1990; Olson & Fazio, 2009).

Experiment 3 went on to further clarify the characteristics of the weighting bias by demonstrating the consequences of having a given weighting bias in a novel situation. In essence, the experiment created invalid, culturally-transmitted prejudices that were positive or negative in nature. A more negative weighting bias was related to giving greater weight to initial false negative information transmitted by an earlier participant. Such individuals avoided testing the associated stimuli fully, and thereby maintained their initial negative evaluations of them. A more positive weighting bias was related to giving less weight to the initial negative information, approaching and exploring the associated stimuli more, and thereby overcoming the initial false information. Because they encouraged fuller testing, invalid positive prejudices were overcome regardless of individuals' valence weighting bias.

Implications

In general, the findings here demonstrate the fruitful application of the behavioral paradigm used in these studies, BeanFest, in measuring an individual's fundamental weighting of resemblance to a positive versus resemblance to a negative when engaged in attitude generalization. In this, as well as in other research, it has been demonstrated that the weighting bias measured in BeanFest is related to a judgments across a variety of domains. As individuals begin BeanFest without prior experience, BeanFest itself represents a relatively pure measure of

how individuals weight positives versus negatives. Due to this measurement approach, as well as the cross-domain nature of valence in general, the weighting bias represents a fundamental individual difference that can influence individuals' behavior across a number of different domains. Indeed, judgments of beans as positive or negative have very little in common with judging ambiguous social situations (e.g., threat assessment, as in Pietri et al. (2013), or risk-taking behavior, as in the current work as well as Pietri et al.) other than the extent to which individuals tend to give weight to positives versus negatives.

When a self-report measure of the weighting bias was administered in Experiment 3, it replicated previous findings showing no correlation with the weighting bias as measured via BeanFest (Pietri et al., 2013). These null findings indicate that the performance-based measure of the weighting bias is non-redundant with what could be obtained through self-report. They also suggest that individuals may have difficulty introspecting and estimating their valence weighting tendencies. As argued in past work (Pietri et al., 2013), it is likely that real-world scenarios do not lend themselves to accurate self-knowledge of whether one tends to give more weight to positives and negatives as valence tends to covary naturally with many other variables, including distinctiveness and diagnosticity (Skowronski & Carlston, 1989).

By demonstrating the weighting of positive versus negative as an initial starting point for individuals when making a judgment or decision, the current research also points to a fundamental way in which individuals navigate their environments and the world in general. Upon entering an environment that requires decision making, individuals will make use of the cues available to them, including attitudes, beliefs, or expectations that might be activated in the situation. In many cases, however, their prior evaluative knowledge may provide little or no guidance regarding this decision process; for example, this occurs when they have little prior

experience or a lack of substantive associations or theories. In situations like these, individuals will need to engage in at least some form of consideration – whether it be more or less quick – to make a decision. Individuals appear to begin such a decision-making process in a somewhat systematic manner by integrating and weighting the overall known positive and negative features of the stimuli. Indeed, individuals’ idiosyncratic tendencies toward valence weighting appear to provide a starting point from which they begin to consider the potential tradeoffs involved and arrive relatively quickly at an overall, gestalt assessment of the situation. After this initial valence weighting, individuals may then deliberate and further reflect on this initial judgment, provided they have both the motivation and opportunity to do so.

The results from Experiment 3 also elucidate the longer-term consequences of having a given weighting bias by showing how such a bias affects approach behavior and attitude change versus maintenance. As argued in the introduction and demonstrated in the final experiment, the weighting bias appears to facilitate attitude formation in at least two different but related ways. First, it seems likely that those with a negative weighting bias had an initial attitude toward the “bad” balloons that was more negative than individuals with a positive weighting bias. After all, it was those with a negative weighting bias who subsequently decided not to test the capacity of these “bad” balloons while those with a positive weighting bias did so. Second, and related to this first point, through this testing or lack thereof, participants’ initial negative attitudes were either overcome or maintained. Those with a negative weighting bias maintained a negative evaluation of the “bad” balloons by inflating them less, whereas those with a positive weighting bias overcame this initial negative information by inflating them more extensively.

Through these two routes, individuals could leave an initially novel environment with entirely different perceptions of how threatening and risky it would be to interact subsequently

with the stimuli, despite the fact that they received the exact same initial information. Beyond reflections about the stimuli they experienced, it also seems likely that these differing perceptions would affect individuals' willingness to re-enter and interact with this environment in the future. It is possible, then, that by leaving an environment with an inaccurate, negative evaluation, those with a negative weighting bias would be depriving themselves of future opportunities for positive outcomes.

In this way, the weighting bias also may relate to recent research concerning what has been termed a "dispositional" attitude. Hepler and Albarracin (2013) have posited that individuals are characterized by a dispositional attitude in the sense that some are more likely to hold positive than negative attitudes, or vice versa. The measure used to assess a dispositional attitude lists 16 attitude objects (e.g., camping, Japan, and rugby) and asks individuals to simply report their attitudes toward them. The more positive their attitudes are on average toward these 16 objects, the more positive the dispositional attitude that is presumed to characterize them. Based on evidence in the current paper, individuals with a more negative valence weighting bias may be more likely to form a negative attitude toward novel objects that are characterized by both positive and negative features. The current findings also suggest that individuals with a more negative weighting bias would be less likely to approach an object toward which they have some pre-existing negative associations, therefore not fully test or explore that object, and thereby maintain their initially negative attitudes toward the object even if they were to be objectively invalid. These general tendencies toward negativity might then be evident on a scale measuring attitudes across multiple objects. Such a link remains to be tested in future research. Its viability likely depends on the currently unknown mechanisms that might underlie scores on Hepler and Albarracin's Dispositional Attitude Measure. A dispositional attitude may stem from

multiple developmental forces, instead of, or in addition to, the particular valence weighting mechanism upon which our more process-oriented individual difference measure focuses.

Although applying the results from valence weighting in novel beans to complex real-world scenarios should be done with caution, it seems plausible to argue that those with a more negative weighting bias may be particularly susceptible to negative stereotypes and prejudices in everyday life. Experiment 3 was framed to participants as a study of how cultural knowledge is passed from generation to generation. Just as in real-world scenarios, participants were given advice from others in the ostensible hope that it might allow them to better navigate their environment by approaching positive stimuli and avoiding negative stimuli. Given the results of Experiment 3, it appears that individuals with a negative weighting bias may give negative culturally-transmitted stereotypes greater credence and develop negative attitudes in line with these stereotypes. Any such negativity is likely to affect their subsequent approach/avoidance behavior. Indeed, research has shown that automatically-activated racial attitudes relate to individuals' willingness to engage with situations involving interracial interactions (Towles-Schwen & Fazio, 2003), that anticipated anxiety about interacting with an African American partner predicts whether individuals will interact with that partner subsequently (Plant & Devine, 2003), and that prejudice has a causal effect on the avoidance of future contact with minority groups which then leads to greater stability of the prejudice (Binder et al., 2009). Given this evidence, once they have developed these initial negative attitudes, those with a negative weighting bias may be all the more likely to maintain those attitudes as they would tend to avoid contact with the stigmatized group in the future.

The final irony of our research is that the culturally-transmitted information in Experiment 3 was actually inaccurate. Past research has demonstrated that, in general, invalid

positive information tends to lead individuals to approach objects and learn their true value. Invalid negative information leads individuals to avoid and to therefore maintain their initial negative attitudes (Fazio et al., 2004). The current research goes a step further by demonstrating that there are certain individuals for whom this is especially true. We have shown that as an individual difference, a positive weighting bias can lead to approach behavior and therefore guide individuals to an accurate understanding of an environment. More insidiously, a negative weighting bias can lead individuals to be more susceptible to invalid negative information and, ironically, see the world as much more threatening than it actually is.

Acknowledgements

We would like to thank Eva Pietri, Alison Young, Pete Zunick, Chris Jones, the Social Cognition Research Group, and the Group for Attitudes and Persuasion for their help and feedback at different points of this research project. We would also like to thank our undergraduate research associates Jordan Tkach, Aaron Hatchett, and Allen Aston for their various contribution to these experiments.

References

- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323-370.
- Binder, J., Zagefka, H., Brown, R., Funke, F., Kessler, T., Mummendey, A., Maquil, A., et al. (2009). Does contact reduce prejudice or does prejudice reduce contact? A longitudinal test of the contact hypothesis among majority and minority groups in three European countries. *Journal of Personality and Social Psychology*, 96(4), 843–856.
- 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(1), 3-25.
- Chaiken, S. & Trope, Y. (Eds.). (1999). *Dual-process theories in social psychology*. New York, NY: Guilford Press.
- Elliot, A. J., & Thrash, T. M. (2010). Approach and avoidance temperament as basic dimensions of personality. *Journal of Personality*, 78(3), 865–906.
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. *Advances in experimental social psychology* (pp. 75-109). New York: Academic Press.
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, 87(3), 293-311.
- Gray, J. A. (1987). *The psychology of fear and stress*. New York: Cambridge University Press.

- Hepler, J., & Albarracín, D. (2013). Attitudes without objects: Evidence for a dispositional attitude, its Measurement, and its consequences. *Journal of Personality and Social Psychology*.
- Higgins, E. T. (1997). Beyond pleasure and pain. *American Psychologist*, 52, 1280-1300.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-291.
- Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., Strong, D. R., et al. (2002). Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, 8(2), 75-84.
- Lejuez, C. W., Aklin, W. M., Jones, H. A., Richards, J. R., Strong, D. R., Kahler, C. W., & Read, J. P. (2003). The Balloon Analogue Risk Task (BART) differentiates smokers from nonsmokers. *Experimental and Clinical Psychopharmacology*, 11, 26–33.
- Lejuez, C. W., Simmons, B., Aklin, W. M., Daughters, S. B., & Dvir, S. (2004). Risk-taking propensity and risky sexual behavior of individuals in residential substance use treatment. *Addictive Behaviors*, 29, 1643–1647.
- Murray, S. L., Holmes, J. G., Collins, N. L., (2006). Optimizing assurance: The risk regulation system in relationships. *Psychological Bulletin*, 5, 641-666.
- Olson, M. A., & Fazio, R. H. (2009). Implicit and explicit measures of attitudes: The perspective of the MODE model. *Attitudes: Insights from the new implicit measures* (pp. 19-63). New York, NY: Psychology Press.
- Pietri, E. S., Fazio, R. H., & Shook, N. J. (2012). Valence weighting as a predictor of emotional reactivity to a stressful situation. *Journal of Social and Clinical Psychology*, 31, 746-777.

- Pietri, E. S., Fazio, R. H., & Shook, N. J. (2013). Weighting positive versus negative: The fundamental nature of valence asymmetry. *Journal of Personality, 81*, 196-208.
- Pacini, R., & Epstein, S. (1999). The relation of rational and experiential information processing styles to personality, basic beliefs, and the ratio-bias phenomenon. *Journal of Personality and Social Psychology, 76*(6), 972-987.
- Plant, E. A., & Devine, P. G. (2003). The antecedents and implications of interracial anxiety. *Personality and Social Psychology Bulletin, 29*, 790–801.
- Preacher, K. J., & Hayes, A. F. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, & Computers, 36*(4), 717–731.
- Rozin, P., & Royzman, E. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review, 5*, 296-320.
- Shook, N. J., Fazio, R. H., & Richard Eiser, J. (2007). Attitude generalization: Similarity, valence, and extremity. *Journal of Experimental Social Psychology, 43*(4), 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.
- Towles-Schwen, T., & Fazio, R. H. (2003). Choosing social situations: The relation between automatically activated racial attitudes and anticipated comfort interacting with African Americans. *Personality and Social Psychology Bulletin, 29*, 170–182.

Figure Captions

Figure 1. Regression lines relating the weighting bias to proportion of donuts approached in DonutFest as a function of condition. Values on the x-axis range from the minimum to the maximum scores within the sample.

Figure 2. Regression lines relating the weighting bias to avoidance behavior in DonutFest as a function of condition. Values on the x-axis range from the minimum to the maximum scores within the sample.

Figure 3. Regression lines relating the weighting bias to balloon pumps in the BART as a function of condition. Values on the x-axis range from the minimum to the maximum scores within the sample.

Figure 4. Graph of three separate linear regression equations relating individuals' weighting bias with their pumping behavior within the BART. Values on the x-axis range from the minimum to the maximum scores within the sample.

Figure 5. Graph of three separate linear regression equations relating individuals' weighting bias with their post-task evaluations of each balloon type. Values on the x-axis range from the minimum to the maximum scores within the sample.

Figure 6. Mediation model of the weighting bias' effect on final attitudes through individuals' pumping behavior. Direct relationship indicated in parentheses. Note: * $p \leq .05$; ** $p < .01$

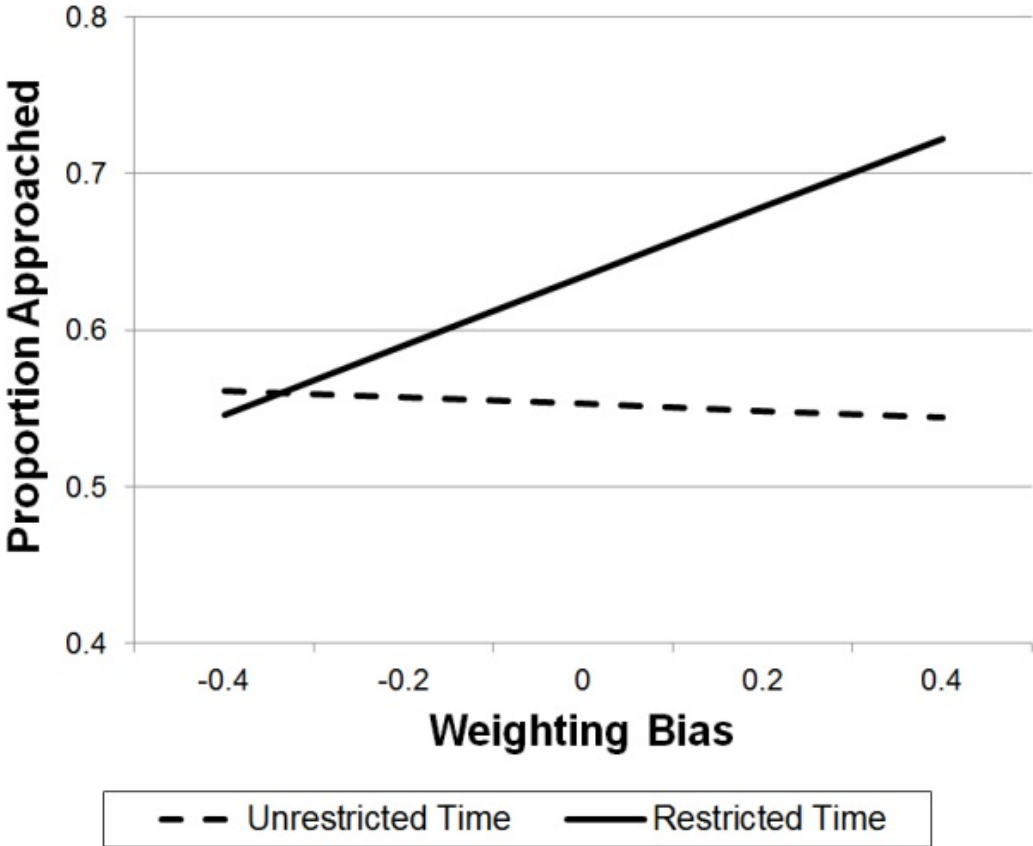


Figure 1

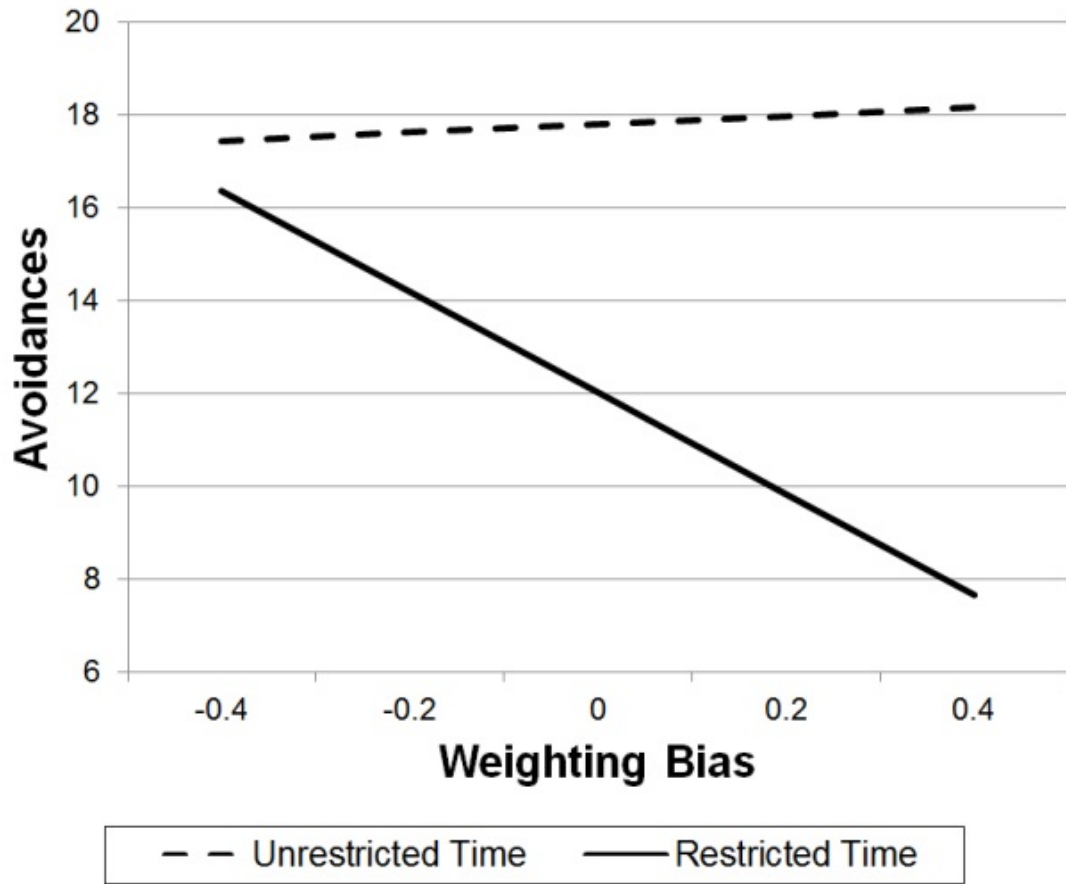


Figure 2

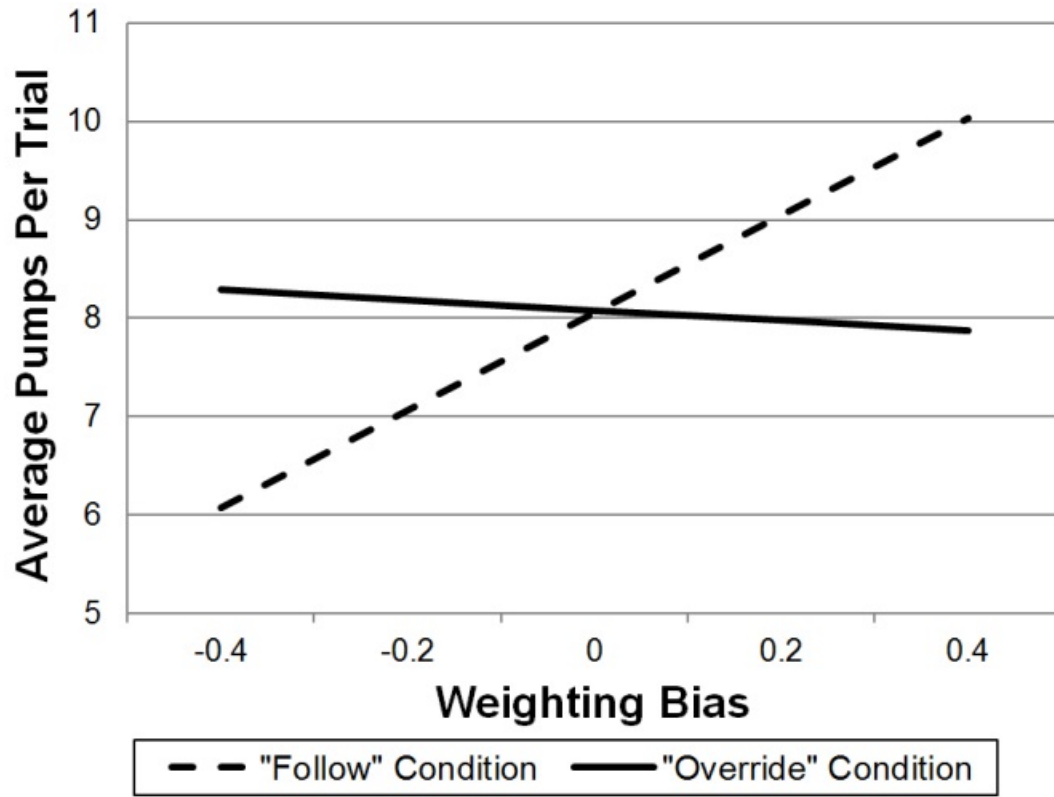


Figure 3

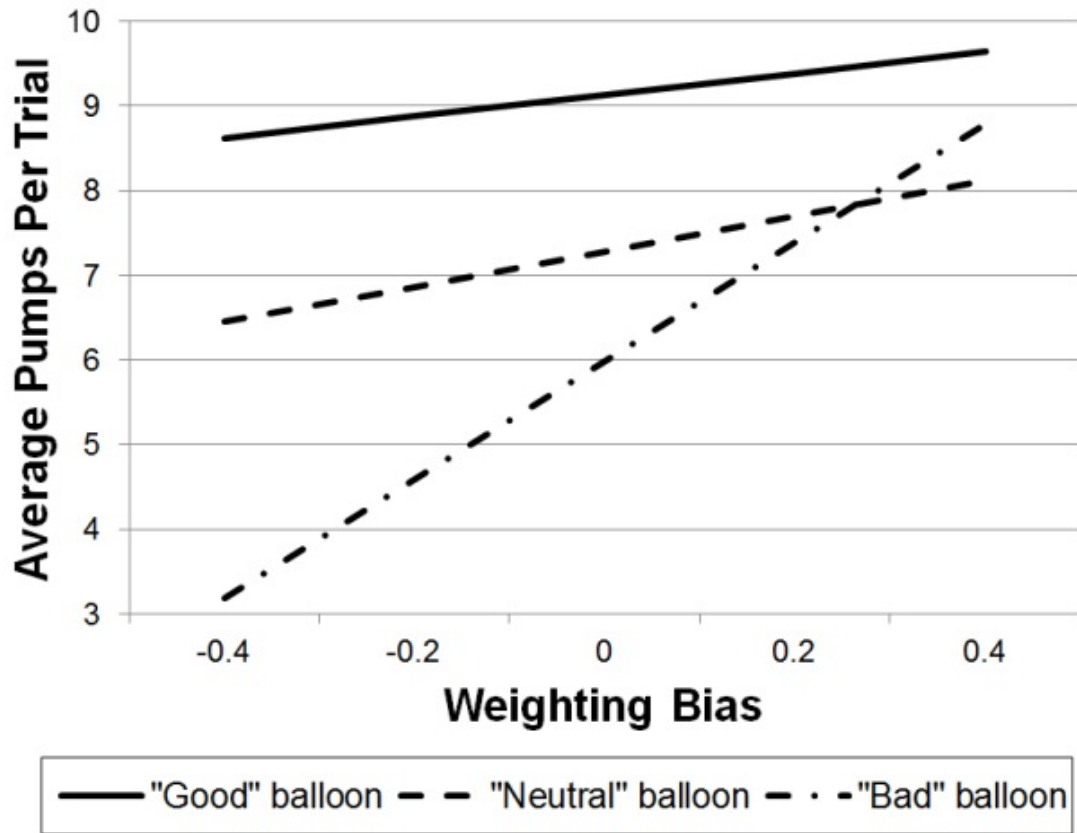


Figure 4

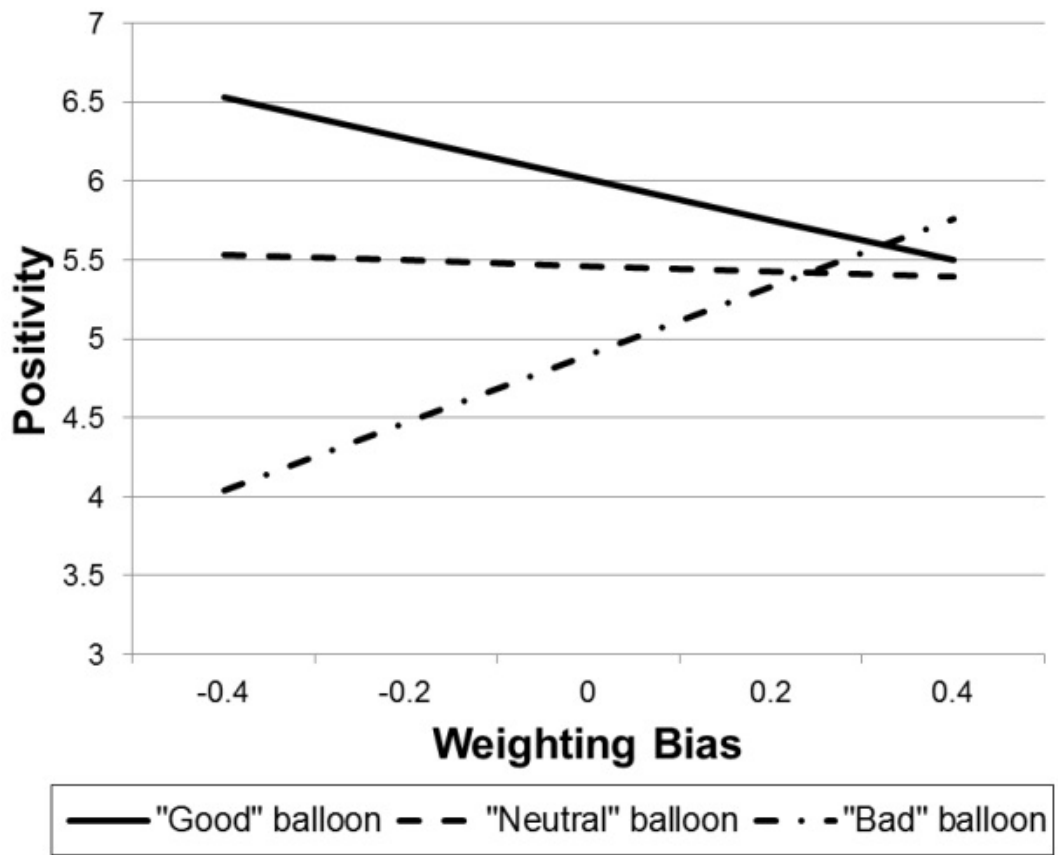


Figure 5

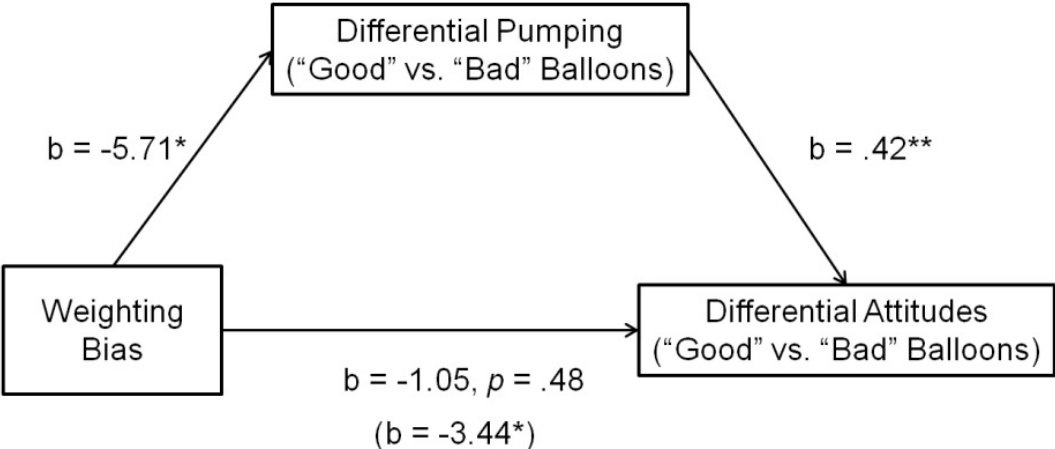


Figure 6