

STICKY FRAMES AND WHAT'S IN A NAME: FRAMES STICK TO PARTICULAR OBJECTS

Y. Andre Wang
University of Toronto

Melisse C. Liwag
Palo Alto University

Katherine Weltzien, Trevor Crowell, and Alison Ledgerwood
University of California, Davis

A growing literature on sequential framing effects has documented a negativity bias: In many contexts, attitudes change less when framing switches from negative-to-positive (vs. positive-to-negative). However, it is unclear whether this negativity bias sticks to one attitude object or generalizes beyond it. Novel paradigms in two experiments yielded strong evidence for the first possibility and tentative evidence for the second: Switching to a different object (vs. same object) across time points reduced the negativity bias in reframing. In contrast, superficially rebranding an object (just changing its name) did not reduce negativity bias. The experiments also provide the first evidence that positive frames are somewhat sticky: A positive initial frame somewhat attenuated the impact of a negative subsequent frame on attitudes. The findings are consistent with the possibility that once an object is framed negatively or positively, that conceptualization sticks in the mind and resists subsequent reframing—especially for negative frames.

Keywords: framing, rebranding, labeling, negativity bias, construal, COAST, functional fixedness

Research across a broad range of disciplines and contexts suggests that describing or framing the same information in either positive or negative terms can dramatically influence people's attitudes and decisions (Krishnamurthy et al., 2001; Levin & Gaeth, 1988; Mao et al., 2021; Quattrone & Tversky, 2004). When information is

This article was accepted under the editorship of Jeffrey W. Sherman.

Address correspondence to Y. Andre Wang, Department of Psychology, University of Toronto, 1265 Military Trail, Toronto, ON M1C 1A4, Canada; or Alison Ledgerwood, Department of Psychology, University of California, Davis, 1 Shields Ave., Davis, CA 95616. E-mail: yilinandre.wang@utoronto.ca or aledgerwood@ucdavis.edu

© 2023 Guilford Publications, Inc.

framed positively (e.g., a program has a 70% success rate), people tend to evaluate it more favorably than when the same information is framed negatively (e.g., the program has a 30% failure rate). This well studied effect is part of a large body of evidence highlighting how people attend to and learn about valenced information, which presumably helps people maximize rewards and avoid punishments when navigating physical and social environments (Chen & Bargh, 1999; Markus & Zajonc, 1985; Pietri et al., 2013a; Rozin & Royzman, 2001).

However, assessing the influence of frame valence on evaluation at a single time point does not accurately reflect most naturalistic learning experiences, which frequently occur through repeated interactions with a stimulus. For example, in everyday life, people might see one news headline that describes a policy as having a 70% success rate, and then see another headline that describes that same policy as having a 30% failure rate. Given that sequencing has proven crucial for understanding psychological processes in a variety of domains (e.g., Asch, 1946; Fang et al., 2018; Gawronski et al., 2010; Hogarth & Einhorn, 1992; Ledgerwood et al., 2011), a full account of how people process valenced information must likewise attend to how that valenced information is sequenced. Thus, recent research has moved beyond examining the effect of a single frame encountered in isolation to studying how people respond to sequentially encountered frames (e.g., Ledgerwood & Boydstun, 2014; Sparks & Ledgerwood, 2019).

In a typical sequential framing paradigm, participants read about an object (e.g., a cognitive training regimen, a surgical procedure, a governor's job record) framed in either positive terms (e.g., it has a 40% success rate) or negative terms (e.g., it has a 60% failure rate). Participants rate their attitude toward the object at Time 1, after encountering just the initial frame. Next, they see the same object reframed in the opposing way (e.g., from success rate to failure rate, or from failure rate to success rate) and rate their attitude again. Whereas the prior literature on single-shot framing effects would have suggested that people's preferences would simply follow the current frame, studies on reframing effects suggest that in many contexts, people's attitudes change less when framing switches from negative to positive, compared to when it switches from positive to negative (Boydston et al., 2019; Itani et al., 2023; Ledgerwood & Boydston, 2014; Sparks & Ledgerwood, 2019). This negativity bias in reframing replicates across multiple cultures and across ages, with the size attenuating as age increases (Itani et al., 2023; Sparks & Ledgerwood, 2019). Researchers have theorized that the negativity bias in reframing reflects a general and presumably adaptive tendency for people to prioritize safety and attend to potential threats in their environment until motivational priorities shift in older age (Sparks & Ledgerwood, 2017, 2019; see also Rozin & Royzman, 2001).

TESTING THE BOUNDARY CONDITIONS OF REFRAMING EFFECTS

Despite considerable evidence for the replicability and generalizability of a negativity bias in reframing, the boundary conditions demarcating the breadth or specificity of this effect are less well understood. In particular, it is unclear whether the negativity bias in reframing (1) specifically affects only evaluations of the initially

framed attitude object, (2) generalizes or carries over from one object to another, or (3) does both.

By better circumscribing the specificity or breadth of the negativity bias in reframing, we can lay an important foundation for future theorizing about this bias. For example, some perspectives place priority on a person's subjective construal of an attitude object (Asch, 1940), suggesting that once a person construes a specific object in a particular way, it can be difficult to reconceptualize that object in a different way (Adamson, 1952; Duncker, 1945; Higgins & Chaires, 1980). Applied to the current context, this perspective suggests that the negativity bias in reframing will be smaller (or disappear entirely) if the attitude object encountered at Time 2 is different from the attitude object encountered at Time 1. Other perspectives assume that valenced stimuli (like an initially encountered frame) can have more diffuse effects, suggesting that the positive or negative affective tone elicited by an initially encountered frame can carry over and affect subsequently encountered and even irrelevant objects (Fazio et al., 1986; Levin et al., 1998; Payne & Lundberg, 2014). In this case, the negativity bias in reframing should persist, even if the attitude object changes across time points. Indeed, across the many forms of negativity bias that have been studied in a wide variety of domains, mechanisms vary in breadth from those that are specific to a particular object to those that carry over or transfer from one object to another (e.g., Ito et al., 1998; Klein & O'Brien, 2016; Ohira et al., 1998; Pietri et al., 2013b; see Rozin & Royzman, 2001, for a review). Clarifying the boundary conditions of reframing effects may therefore provide a key first step in assessing the plausibility of various potential mechanisms and in guiding our attention to relevant theorizing across these disparate literatures.

Moreover, understanding the extent to which negative frames can have lingering effects on different objects (vs. only the initially framed object) could have important applied implications. A political candidate might make different framing choices when debating if an initial negative frame will affect only the framed issue versus a second, subsequently described issue. Someone making recommendations for organizational policy might want to know if highlighting the downsides of one policy alternative would dull the impact of positively framing a second policy alternative. Thus, for both theoretical and applied reasons, it seems important to test the extent to which reframing effects can carry over from one attitude object to another.

THE PRESENT RESEARCH

To begin characterizing the specificity or breadth of reframing effects, we first needed to create two versions of a reframing paradigm that would allow us to distinguish between object-specific and object-general effects. An object-specific prediction suggests that a negativity bias in reframing will specifically follow (and not generalize beyond) the particular object that is framed at Time 1. In contrast, an object-general prediction suggests that the negativity bias in reframing will generalize to a different object encountered at Time 2. We therefore manipulated

whether the initially framed object at Time 1 remained the same at Time 2 (as in past research) or switched to a new object at Time 2.

In Study 1, participants learned about a program that would lead to a certain percentage of teaching jobs being saved (positive frame first) or lost (negative frame first) and rated their attitudes toward that program at Time 1. We then varied what happened at Time 2. In the same-program condition, participants saw the same teaching jobs program reframed in negative or positive terms and then rerated their attitudes toward the same program at Time 2. In the different-program condition, participants instead saw a health care jobs program framed in negative or positive terms and then rated their attitudes toward this different program at Time 2. By comparing the same-program and different-program conditions, we could assess the extent to which the negativity bias in reframing was specific to the initially framed object or generalized across objects.

In Study 2, we sought to replicate and extend our findings to probe whether even relatively superficial changes to an object would be sufficient to reduce the negativity bias in reframing. We therefore created a third condition in which we kept the object the same across time points but “rebranded” it at Time 2 by giving it a new name and claiming it was a “different version” of the same program. By comparing the rebranded condition to the same-object and different-object condition, we could more precisely assess the specificity of sequential framing effects and shed some experimental light on whether and how rebranding might be an effective strategy for negatively framed policies.

Materials, datasets, and data analysis scripts for all studies are available at <https://osf.io/6gec8/> (Wang et al., 2022).

STUDY 1

We designed Study 1 to examine the two competing possibilities for the negativity bias in sequential framing effects. Following paradigms of past research on this topic (e.g., Ledgerwood & Boydstun, 2014; Sparks & Ledgerwood, 2017, 2019), participants evaluated programs addressing budget cuts that were initially framed either positively (i.e., in terms of jobs saved) or negatively (i.e., in terms of jobs lost) at Time 1 and then subsequently reframed at Time 2. After each frame, participants indicated their attitude toward the program. Participants were randomly assigned to one of the two framing orders: an initial positive frame followed by a subsequent negative reframe (*positive-to-negative* condition), or an initial negative frame followed by a subsequent positive reframe (*negative-to-positive* condition). Importantly, we also manipulated whether the program was the same or different at Time 1 versus Time 2: Some participants evaluated the same program at both time points (*same-program* condition), whereas other participants evaluated a different program at Time 1 versus Time 2 (*different-program* condition). If the negativity bias in sequential framing is object-general, then the extent of negativity bias should not depend on whether the program is the same or different from Time 1 to Time 2. In contrast, if the bias is object-specific, then the extent of negativity bias should diminish when the program is different (vs. the same).

METHOD

Participants and Power. Three hundred and thirteen participants who had not participated in a previous study on reframing effects completed the experiment online through Amazon's MTurk platform.¹ They were randomly assigned to one cell of a 2 (framing order: positive-to-negative vs. negative-to-positive) \times 2 (evaluated program: same vs. different) between-subjects design. We based our target sample size on prior research on sequential framing effects, which has found a meta-analytic effect size of $d = 0.54$ for negativity bias in the loss domain (Studies 1, 2, and the study reported in Footnote 4; Ledgerwood & Boydstun, 2014). Power analysis conducted in G*Power (Faul et al., 2007) suggested that a cell size of $n = 60$ would allow us to detect a simple main effect of this size in the same program condition at over 80% power, so we set a total target sample size of $N = 240$. Based on pilot testing exploring how carefully participants attended to the same versus different program manipulation, we anticipated an exclusion rate of about 20%. We therefore planned to collect data from 300 participants (we followed this stopping rule but the total number of completed surveys in Qualtrics ended up being slightly higher).

We set and recorded key researcher decisions on a time-stamped internal form, following our lab's standard protocol for documenting every study on a shared lab drive to support our ability to distinguish between data-independent and data-dependent decisions and limit unintended Type I error inflation (see da Silva Frost & Ledgerwood, 2020, for more on different forms of preregistration and the different goals they can serve). We set and recorded the following a priori exclusion criteria: We would exclude participants who (1) failed any attention check on the policies they were assigned to learn about, (2) failed to respond to two or more items of the key dependent variable at either time points, and/or (3) indicated suspicion of the hypothesis or knowledge of framing effects. The numbers of participants who met each of these exclusion criteria were 20, 1, and 13, respectively (some participants met multiple exclusion criteria).² The final sample size was 280 (108 men, 170 women, 2 did not provide gender information; $M_{\text{age}} = 38.3$, $SD_{\text{age}} = 12.0$; 6% Asian, 12% Black/African American, 2% multiracial, 1% Native American or Alaska Native, 79% White).³

Materials and Procedure. Participants were asked to imagine that "an important state is addressing statewide budget cuts, which are expected to affect jobs in

1. In this study and in Study 2, we followed the lab's standard operating procedures and counted participants who completed the final attention check as completing the study.

2. The numbers of participants who met each of these exclusion criteria by condition were, respectively: 4, 1, 1 (negative-to-positive, same program); 7, 0, 10 (negative-to-positive, different program); 2, 0, 0 (positive-to-negative, same program); 7, 0, 2 (positive-to-negative, different program).

3. Due to a survey design oversight, we did not directly ask participants about their race and ethnicity in our studies. For both this study and Study 2 below, the percentages of participants' race and ethnicity were estimated from meta-survey data on the MTurk workers who completed the study and might not reflect the exact percentages for participants in the final sample.

teaching, healthcare, manufacturing, and information technology,” and that “[s]everal programs have been proposed to address possible job losses” in each of those areas. They were then introduced to “Program T,” a teaching jobs program proposed to address the teaching jobs that were at stake (Time 1). Participants learned about experts’ consensus on the outcome of this program, which were framed either positively (“40% of these teaching jobs will be saved”) or negatively (“60% of these teaching jobs will be lost”). To ensure that participants were attending to the evaluated program manipulation, we bolded the domain of the program (e.g., “teaching jobs program”) throughout the study (the bolded domain information appeared six times in total). In addition, we included an attention check at each time point just below the program information that read: “For verification, please indicate which program is described above” and asked participants to select between a correct option (“Program T”) and an incorrect option (“Program H”). Participants then rated their attitudes toward Program T by moving sliders along three 100-point scales anchored at the endpoints (*completely oppose Program T to completely favor Program T, not at all positive to extremely positive, and not at all negative to extremely negative* [reverse-coded]). These items were averaged to form an index of participants’ attitude at Time 1 ($\alpha = .97$).

Next, participants saw additional information. In the *same-program* condition, the “additional information” simply reframed the prior information about Program T with the opposite frame valence. For example, participants who read at Time 1 that 40% of teaching jobs would be saved under Program T (positive frame) read “additional information” at Time 2 that described the outcome of the program as “60% of these teaching jobs will be lost” (negative frame). In the *different-program* condition, the “additional information” described a different program, Program H, proposed to address the health care jobs that were at stake. Participants in this condition then saw the reframed information for Program H. For example, participants who read at Time 1 that 40% of teaching jobs would be saved under Program T (positive frame) read at Time 2 that 60% of the health care jobs would be lost under Program H (negative frame). The full list of frames at Time 1 and Time 2 for each condition is reported in Table 1. After completing the Time 2 attention check on which program had just been described, participants then indicated their attitudes toward the program at Time 2 (Program T for the same-program condition; Program H for the different-program condition) using the same three scales, which were again averaged to form an index of attitude at Time 2 ($\alpha = .95$).⁴ Participants then completed a final attention check that asked them to identify which program(s) they had rated on the past few screens out of a set of options that included Program T, Program H, and multiple distractor items. Last, participants provided their demographic information and completed a funnel debrief that probed for suspicion of the hypothesis

4. Following past studies in this program of research (e.g., Sparks & Ledgerwood, 2017), we collected data on a series of exploratory measures, including participants’ emotional states, the extent they deliberated on the questions, the amount of attention they paid to the questions, information processing ease, and general attention to the survey. These measures appeared after the DVs as part of a standard inventory of exploratory measures. We did not plan to analyze these measures because they were not central to our key research questions, and they are not discussed further. The full list and wording of these measures are available in the study materials on OSF.

TABLE 1. Frames for Each Condition in Study 1

Framing order	Evaluated program	Time	
		Time 1	Time 2
Positive-to-negative	Same	Program T: "40% of these teaching jobs will be saved"	Program T: "60% of these teaching jobs will be lost"
	Different	Program T: "40% of these teaching jobs will be saved"	Program H: "60% of the healthcare jobs will be lost"
Negative-to-positive	Same	Program T: "60% of these teaching jobs will be lost"	Program T: "40% of these teaching jobs will be saved"
	Different	Program T: "60% of these teaching jobs will be lost"	Program H: "40% of the healthcare jobs will be saved"

and knowledge of framing effects (e.g., "Do you think that anything you did in one part of the study could have influenced you at another part? If so, how?"), after which they were debriefed.

RESULTS

Following our internal pre-analysis plan, we conducted a 2 (framing order: positive-to-negative vs. negative-to-positive) \times 2 (evaluated program: same vs. different) between-subjects analysis of variance (ANOVA) on the extent of attitude change that participants displayed in the direction of the Time 2 frame (i.e., the amount of attitude change that was shifted away from the Time 1 frame and toward the Time 2 frame). As predicted, there was a main effect of framing order, $F(1, 276) = 16.35, p < .001, \eta_p^2 = .06$, 90% confidence interval (CI) [.02, .10], suggesting that attitude change was significantly smaller when a negative frame was followed by a positive reframe ($M = 14.50, SD = 29.39$) than when a positive frame was followed by a negative reframe ($M = 29.53, SD = 32.69$).⁵

More central to the purpose of our study, the framing order \times evaluated program interaction was significant, $F(1, 276) = 4.66, p = .032, \eta_p^2 = .02$, 90% CI [.001, .05], suggesting that the effect of framing order differed by whether the evaluated program was the same (vs. different) across the two time points (see Figure 1). Consistent with past research, in the same-program condition, attitudes changed significantly less when the frame switched from negative to positive ($M = 2.18, SD = 13.76$), compared to when the frame switched from positive to negative ($M = 24.07, SD = 32.23$), $t(276) = 4.57, p < .001, d = 0.74$, 95% CI [0.42, 1.07], with a medium-large effect size. In contrast, in the different-program condition, attitude change was statistically comparable when the framing switched from negative to positive ($M = 29.15, SD = 35.79$), compared to when it switched from positive to

5. We did not plan to test the main effect of program because it was not directly relevant to the purpose of our study. Note that the historically common practice of testing and reporting all effects in a multiway ANOVA can dramatically inflate Type I error rates (Cramer et al., 2016).




FIGURE 1. Attitude change by frame valence order and evaluated program in Study 1. Error bars depict ± 1 standard errors; dots depict jittered individual data points.

negative ($M = 35.81$, $SD = 32.32$), $t(276) = 1.28$, $p = .200$, $d = 0.23$, 95% CI $[-0.12, 0.57]$, with a small effect size in the same direction. Figure 2 is presented for readers who prefer to see attitudes plotted separately at each time point.

Our study design also offered an opportunity to explore reframing effects separately for positive-to-negative and negative-to-positive reframing. We explored (in an unplanned analysis) pairwise comparisons for the effect of program in the positive-to-negative and negative-to-positive conditions separately. In the positive-to-negative condition, attitudes changed significantly less when the program was the same ($M = 24.07$, $SD = 32.23$), compared to when it was different ($M = 35.81$, $SD = 32.32$), $t(276) = 2.37$, $p = .019$, $d = 0.40$, 95% CI $[0.07, 0.73]$, with a small-to-medium effect size. In the negative-to-positive condition, attitudes changed significantly less when the program was the same ($M = 2.18$, $SD = 13.76$), compared to when it was different ($M = 29.15$, $SD = 35.79$), $t(276) = 5.36$, $p < .001$, $d = 0.92$, 95% CI $[0.57, 1.26]$, with a large effect size.

DISCUSSION

In Study 1, we found that the size of the negativity bias in reframing diminished when the program was different (vs. the same) from Time 1 to Time 2. Whereas we replicated the typical negativity bias observed in past research in the same-program




FIGURE 2. Attitude by framing order, time point, and evaluated program in Study 1. Error bars depict ± 1 standard errors; dots depict jittered individual data points.

condition, with a medium-to-large effect size, the bias decreased to a nonsignificant and small effect size in the different-program condition. These findings are consistent with the notion that sequential framing effects are at least largely object-specific. If they were completely object-general, then we should have observed a similar pattern of results regardless of whether participants saw the same program or a different program from one time point to the next.

Notably, the observed pattern of results also helps rule out a possible alternative account for reframing effects, which is that that positive information is simply ineffective, regardless of when it is encountered. Such an account would suggest that the so-called “negativity bias in reframing” observed in past research and in the same-program condition of this study is not actually about negativity or reframing at all, but rather reflects a general tendency for positive (vs. negative) information to be unpersuasive or ineffective. In the absence of a carefully constructed comparison condition, it is impossible to tell whether the pattern of results observed in prior studies on reframing (e.g., *Boydston et al., 2019; Ledgerwood & Boydston, 2014*) reflects stickiness in the initial negative frame or weakness in the subsequently encountered positive frame. Specifically, it could be that attitudes change less in the negative-to-positive (vs. positive-to-negative) reframing conditions because at Time 2, the information contained in the positive (vs. negative) reframe simply has a weaker impact on attitudes. However, if this were in fact the case, a positive frame at Time 2 should be equally ineffective regardless of whether it is framing the same program encountered at Time 1 or a different program. Instead, we see that the effect of the positive Time 2 frame changed depending on whether the framed program was the same or different from one time point to the next.

Finally, it is interesting to note the unplanned (exploratory/data-dependent) simple main effect of program in the positive-to-negative framing condition. Whereas prior research had only been able to offer evidence that negative frames are relatively stickier than positive frames, the design of Study 1 offered a first opportunity to explore the possibility that positive frames may also be sticky. Does an initial positive frame mute the impact of a subsequently encountered negative frame when the framed object is the same (vs. different) across time points? The small-to-medium and significant effect of program (same vs. different) in the positive-to-negative framing condition provides tentative evidence that a positive frame may also stick to a particular object and resist reframing, just to a lesser extent than a negative frame. However, given the data-dependent nature of this test and the relatively high p value ($p = .019$), such an effect would need to be replicated before we could draw any conclusions with much confidence.

STUDY 2

The results of Study 1 provide crucial new evidence for the specificity of the negativity bias in reframing, suggesting that the stickiness of an initial negative frame attaches to and follows a particular object. These initial findings raise intriguing questions about what it takes to unstick a negative frame once it has been applied to a particular object. Can “rebranding” or superficially changing an object operate like the different-object condition in Study 1, releasing an object from an initial negative frame? Or will an initial negative frame follow an object regardless of how it is described?

People who are in the role of marketing an organization or policy may intuitively think that rebranding should work to circumvent an initial negative conceptualization of an object. For example, in October 2021, in the wake of immense legal and public scrutiny, social media giant Facebook unveiled its intention to initiate a comprehensive rebranding effort, renaming itself Meta and changing its logo (Isaac, 2021). According to cofounder and CEO Mark Zuckerberg, the change would represent Facebook’s de-emphasis of its social media platform to the end of more fervently pursuing what he contended to be the next high-tech frontier: the unification of disparate digital spaces into a single augmented reality called the “metaverse” (Isaac, 2021). However, some speculated that the rebrand was largely driven by a desire to distance the company from its fraught history and increasingly negative reputation (e.g., multiple congressional hearings, a lawsuit by the Federal Trade Commission, and a fine of \$5 billion for deception; Paul, 2021).

Thus, people seem to intuitively expect that rebranding a company or a policy will provide a fresh start. This expectation has been documented in numerous case studies examining the use of rebranding strategies to improve public attitudes toward corporations and political or public health terminology (Muzellec & Lambkin, 2006; Smith & Malone, 2003; Sørensen et al., 2020; Witmarsh, 2009). However, these case studies present mixed findings for the effectiveness of rebranding when used to distance a corporation or term from existing negative associations (Muzellec & Lambkin, 2006; Witmarsh, 2009). For example, Muzellec and Lambkin (2006)

contrasted the rebranding of two telecommunication companies in Ireland and found that whereas the rebranding effort of one company drastically improved its brand image, a similar effort of the other company had mixed success and resulted in worse evaluations among some of its customers. Moreover, there is a dearth of experimental work examining rebranding's effectiveness through a social cognitive lens. Still, we can find some indirect support for the potential effectiveness of rebranding in social-cognitive work suggesting that using different labels for the same attitude object (e.g., "African American" vs. "Black" when describing the same social group) can elicit different evaluative responses by activating different associated attributes (Hall et al., 2015; Philogène, 2001). Similarly, activating different exemplars associated with a social category (e.g., a liked vs. disliked politician) can lead to different evaluative responses (Asch, 1940; Sia et al., 1997). However, this research has focused on familiar attitude objects, exemplars, and labels, leaving unanswered the question of how different valences attach to attitude objects in the early stages of attitude formation (see also da Silva Frost et al., 2022). Does relabeling an attitude object create a blank slate, releasing people from the stickiness of an initial negative frame? Or would reframing effects for a relabeled object look more similar to the same object condition in Study 1, suggesting that initially encountered negative frames attach more deeply to the novel attitude object itself, regardless of how it is subsequently labeled?

In Study 2, we sought to replicate and extend Study 1 by adding a new condition to our design: In addition to the *same-program* and *different-program* conditions, we added a third condition in which participants read about a teaching jobs program at Time 1 and then read about a rebranded version of that same program at Time 2 (*rebranded-program* condition). We were interested in whether the pattern of results in the *rebranded* condition would look like the *different-program* condition (suggesting that an initially encountered negative frame attaches very specifically to a particular object described in a particular way) or instead like the *same-program* condition (suggesting that an initially encountered negative frame can attach to and follow an object regardless of how it is described).

METHOD

Participants and Power. Participants ($N = 2,115$) who had not participated in a previous study on reframing effects completed the experiment online through Amazon's MTurk platform. They were randomly assigned to one cell of a 2 (framing order: positive-to-negative vs. negative-to-positive) \times 3 (evaluated program: same vs. different vs. rebranded) between-subjects design. As in Study 1, we set and recorded key researcher decisions on a time-stamped internal form. We set the same a priori exclusion criteria as Study 1 and a target sample size of 720 ($n = 120$ per condition). A sensitivity analysis in G*Power indicates that a cell size of $n = 60$ provides 80% power to detect an effect size of $d = 0.52$ (which is approximately the size of the negativity bias detected in past research; Ledgerwood & Boydston, 2014); using the rules of thumb for powering interactions described in da Silva Frost and Ledgerwood (2020), we can estimate that doubling this cell size

to $n = 120$ per cell ($N = 720$ total) would provide 80% power to detect a knock-out interaction in the same program versus different program conditions (if the design were the same 2×2 as Study 1) and presumably less power to detect a 2×3 interaction in the current design (but rules of thumb for such interactions are still lacking). To balance power considerations and resource constraints, we planned a group sequential analysis that would allow us to stop data collection early if we find clear support for our target effect of framing order \times evaluated program interaction, while preserving Type I error rate at 5% (Lakens, 2014; Ledgerwood et al., 2017). Specifically, we planned two equally spaced analyses after collecting 360 and 720 participants, with alpha boundaries set at .025 and .034, respectively (da Silva Frost & Ledgerwood, 2020, Table 2; these boundaries can also be calculated using the linear spending α^*t^ϕ function in the R package GroupSeq; Pahl, 2022).

Despite our careful sample size planning, the p values of the target framing order \times evaluated program interaction from the two analyses were $p = .306$ and $p = .166$, which were above their respective alpha cutoffs and yet frustratingly uninformative: After investing substantial resources into the study, it was unclear from the pattern of results whether we could conclude that there were really no differences between conditions or whether our study was simply underpowered to detect them. This ambiguous gray area is exactly what an underappreciated analytic approach called COAST (composite open adaptive sequential test; Frick, 1998) was designed to help navigate. The COAST approach, similar to the group sequential analysis, is another optional stopping technique that controls the overall Type I error rate. COAST allows researchers to perform the same statistical test repeatedly (as many times as desired and with any spacing, making it more flexible than the group sequential analysis), as long as the researchers follow COAST's decision rule of when to stop data collection (see also Sagarin et al., 2014). This decision rule states that to maintain a Type I error rate of $\alpha = .05$, researchers should stop data collection and reject the null hypothesis when $p < .010$ (note that this is a stricter threshold of significance than the typical $p < .050$), and that they should stop data collection and not reject the null hypothesis when $p > .360$.

Because we had thus far followed the COAST decision rules (namely, continuing data collection when the p value of the target framing order \times evaluated program interaction fell between the COAST decision thresholds), we were able to continue using COAST to learn more from our data without fear of inflating our Type I error over our desired rate of .05. Therefore, we decided to follow COAST and collect additional data in batches of 120 participants until the p value from the cumulative sample was out of the $[.010 < p < .360]$ range. Doing so allowed us to maximize what we could learn from the substantial amount of data we had already collected while ensuring that the repeated testing did not inflate our Type I error rate (da Silva Frost & Ledgerwood, 2020; Ledgerwood et al., 2017). The p values of the target framing order \times evaluated program interaction of the cumulative sample after each additional batch of participants were .075, .049, .035, .021, .049, .025, .025, .026, .021, and .008, at which point the p value fell outside of the decision range, triggering the end of data collection.

We followed our a priori exclusion criteria. Of the 2,115 participants who completed the study, 195 failed at least one attention check on the policies they were assigned to learn about, 4 failed to respond to two or more items of the key dependent variable at either time points, and 272 indicated suspicion of the hypothesis or knowledge of framing effects (some participants met multiple exclusion criteria), leaving a final sample size of 1,654 (595 men, 1,043 women, 9 reported another gender identity, 7 did not provide gender information; $M_{\text{age}} = 37.7$, $SD_{\text{age}} = 12.6$; 6.6% Asian, 11.4% Black/African American, 2.6% Multiracial, 0.6% Native American or Alaska Native, 78.8% White).⁶

Materials and Procedure. The materials and procedure were similar to Study 1, except for the following changes. We added a third evaluated program condition. In this new, *rebranded-program* condition, participants read about a teaching jobs program (called Program T) at Time 1 that was framed in either negative or positive terms, just as in the other conditions. At Time 2, they read that “a different version of the teaching jobs program (called Program H) has been proposed,” and then (just as in the other conditions) saw the program framed in either positive or negative terms. In all conditions, we used the same three-item index of attitudes at Time 1 ($\alpha = .95$) and Time 2 ($\alpha = .94$).

RESULTS

Following our internal pre-analysis plan, we conducted a 2 (frame valence order: negative-to-positive vs. positive-to-negative) \times 3 (evaluated program: same vs. different vs. rebranded) between-subjects ANOVA on the extent of attitude change participants showed in the direction of the Time 2 frame. There was a significant main effect of frame valence order, $F(1, 1648) = 144.31$, $p < .001$, $\eta_p^2 = .08$, 90% CI [.06, .10], indicating that overall, attitudes changed less when framing switched from negative to positive (vs. positive to negative).

More central to the purpose of our study, the interaction between frame valence order and evaluated program emerged from the COAST procedure on the significant side of the predetermined and ambiguous [$.010 < p < .360$] range, $F(2, 1648) = 4.88$, $p = .008$, $\eta_p^2 = .006$, 90% CI [.0009, .01], indicating that the size of the negativity bias in reframing varied by evaluated program condition (see Figure 3).⁷ Descriptively, the size of the negativity bias was largest in the same-program condition, with a medium-large effect size estimate that was comparable to that observed in Study 1: Attitudes changed substantially less when the frame switched from negative to positive ($M = 6.04$, $SD = 18.07$), compared to when the frame switched from positive to negative ($M = 27.93$, $SD = 28.62$), $t(1648) = 8.87$,

6. The numbers of participants who met each of these exclusion criteria by condition were, respectively: 19, 1, 23 (negative-to-positive, same program); 29, 0, 51 (negative-to-positive, different program); 40, 1, 86 (negative-to-positive, rebranded program); 12, 0, 19 (positive-to-negative, same program); 39, 2, 31 (positive-to-negative, different program); 56, 0, 62 (positive-to-negative, rebranded program).

7. Again, we did not plan to test the main effect of program because it was not directly relevant to the purpose of our study (see Cramer et al., 2016).




FIGURE 3. Attitude change by frame valence order and evaluated program in Study 2. Error bars depict ± 1 standard errors; dots depict jittered individual data points.

$p < .001$, $d = 0.71$, 95% CI [0.55, 0.86]. The size of the negativity bias was smallest in the different-program condition, with a small-to-medium effect size estimate that was slightly larger than that observed in Study 1 (though the confidence intervals largely overlapped): Attitudes changed somewhat less when the frame switched from negative to positive ($M = 21.53$, $SD = 32.80$), compared to when the frame switched from positive to negative ($M = 33.26$, $SD = 36.71$), $t(1648) = 4.44$, $p < .001$, $d = 0.38$, 95% CI [0.21, 0.55]. Finally, in our new, rebranded-program condition, the size of the negativity bias was medium-large and essentially identical to the same-program condition: Attitudes changed substantially less when the frame switched from negative to positive ($M = 16.18$, $SD = 34.96$), compared to when the frame switched from positive to negative ($M = 37.93$, $SD = 33.85$), $t(1648) = 7.60$,

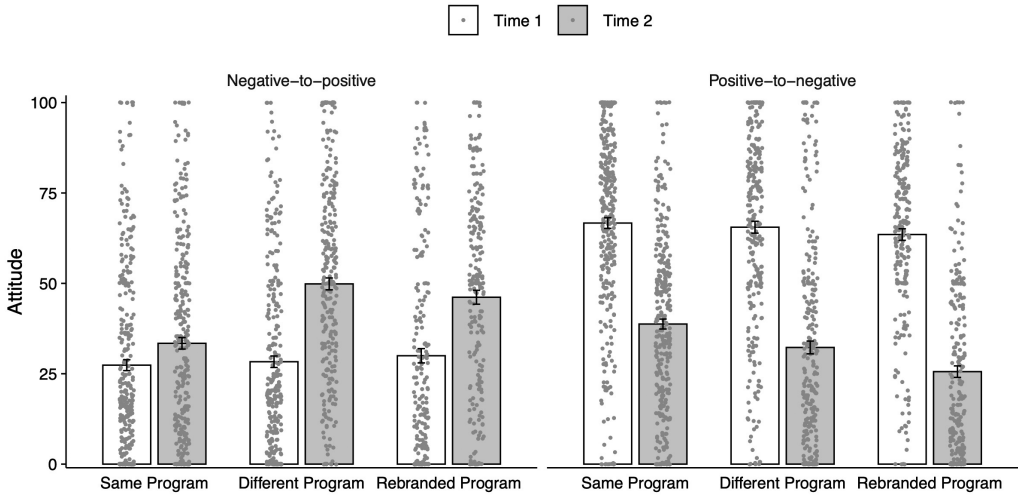


FIGURE 4. Attitude by framing valence order, time point, and evaluated program in Study 2. Error bars depict ± 1 standard errors; dots depict jittered individual data points.

$p < .001$, $d = 0.70$, 95% CI [0.52, 0.88].⁸ Figure 4 is presented for readers who prefer to see attitudes plotted separately at each time point.

To explore whether these effect sizes differed from each other, we conducted (unplanned/exploratory) tests of the 2×2 interactions between each pair of evaluated program conditions. Replicating the 2 (evaluated program: same vs. different) $\times 2$ (frame valence order: negative-to-positive vs. positive-to-negative) interaction from Study 1, the size of the negativity bias in reframing was significantly smaller in the different- (vs. same-) program condition, $F(1, 1180) = 8.71$, $p = .003$, $\eta^2_p = .01$, 90% CI [.001, .02]. Focusing next on the size of the negativity bias in the new rebranded-program condition, it was significantly larger than the different-program condition, $F(1, 1018) = 5.31$, $p = .021$, $\eta^2_p = .005$, 90% CI [.0004, .02], but not the same-program condition, $F(1, 1098) = 0.00$, $p = .967$, $\eta^2_p = .00$, 90% CI [.00, 1.00]. In other words, rebranding the original program was not sufficient to reduce the size of the negativity bias in reframing, whereas switching to a different program did reduce the bias.

Finally, we examined whether the exploratory analyses in Study 1 replicated in Study 2 by comparing just the same- versus different-program conditions within the positive-to-negative and negative-to-positive conditions separately. In the positive-to-negative condition, attitudes changed significantly less when the program was the same ($M = 27.93$, $SD = 28.62$), compared to when it was different ($M = 33.26$, $SD = 36.71$), $t(1180) = 2.20$, $p = .028$, $d = 0.18$, 95% CI [0.02, 0.34], with

8. Note that the statistical tests accompanying these descriptive effect sizes were not planned a priori and should therefore be interpreted with more caution than a planned or data-independent test (da Silva Frost & Ledgerwood, 2020), although we feel fairly confident in them given the very low p values (see McShane et al., 2019).

a small effect size that was about half the size observed in our initial exploratory analysis in Study 1. In the negative-to-positive condition, attitudes changed significantly less when the program was the same ($M = 6.04$, $SD = 18.07$), compared to when it was different ($M = 21.53$, $SD = 32.80$), $t(1180) = 6.31$, $p < .001$, $d = 0.52$, 95% CI [0.36, 0.69], with a medium effect size that was also somewhat smaller than the estimate from the exploratory analysis in Study 1. In other words, the comparisons were again significant in Study 2, suggesting that we can have more confidence that these effects exist, but the estimates were substantially smaller in size, suggesting that our initial exploratory/data-dependent analyses in Study 1 probably overestimated the size of these effects (see Ledgerwood et al., 2017).

DISCUSSION

The results of Study 2 replicate and extend the results observed in Study 1. Once again, the size of the negativity bias in reframing changed depending on evaluated program condition, suggesting that reframing effects are object-specific. We observed very similar effect sizes to Study 1 in the same-object condition, where the negativity bias was medium-to-large, and in the different-object condition, where the negativity bias was small-to-medium.

We also introduced a new, rebranded-program condition in Study 2 to examine whether relatively superficial changes to an object would be sufficient to reduce the negativity bias in reframing, as the popularity of rebranding in the business world and as research on relabeling familiar attitude objects (that are already well connected to both positive and negative associations) might suggest. The answer to this question appears to be a resounding no: The negativity bias in the rebranded-program condition looked virtually identical to the negativity bias in the same-program condition. This pattern of results suggests that an initially encountered negative (vs. positive) frame can more strongly attach to and follow an object, regardless of how that object is labeled.

Interestingly, even though rebranding did not reduce negativity bias in reframing, it did increase attitude change. Compared to the same-program condition, attitude change was generally greater in the rebranded condition: Descriptively, an exploratory/unplanned pairwise comparison (using the pooled error term from the main 2×3 model) estimated this effect as $d = 0.32$, 95% CI [0.20, 0.44]. This pattern of results was not part of a focal, data-independent analysis and should therefore be interpreted tentatively, but it suggests the intriguing possibility that relabeling an attitude object might similarly release people at least a little from an initial positive *or* negative conceptualization, or perhaps make people especially susceptible to any information encountered at Time 2. Thus, it is possible that rebranding might be an effective strategy to enhance the effects of reframing in general (regardless of whether it was initially framed in positive or negative terms), but it does not seem to reduce the negativity bias in reframing.

Finally, in Study 2, we again found a significant (though small) difference between the same- and different-program conditions in the positive-to-negative condition,

allowing us to boost our confidence in what had been a tentative conclusion about this effect in Study 1. We return to this finding in the General Discussion.

GENERAL DISCUSSION

People often encounter the same information sequentially framed in positive and negative terms (e.g., a program described in terms of its success rate or its failure rate). A growing body of research suggests that in many contexts, attitudes change less in response to reframing when information is initially framed in negative (vs. positive) terms (Boydston et al., 2019; Itani et al., 2023; Ledgerwood & Boydston, 2014; Sparks & Ledgerwood, 2019). Across two experiments, we sought to better characterize the boundary conditions of this negativity bias in reframing by designing a paradigm that would allow us to assess the specificity or breadth of these effects. We found strong and consistent evidence that the stickiness of an initial negative frame attaches to the specific attitude object being framed. More specifically, the negativity bias in reframing was reduced when participants encountered a different attitude object across the initial and subsequent frames (e.g., one program framed in negative terms and then a different program framed in positive terms), compared to when participants encountered the same attitude object at both time points (e.g., one program framed in negative terms and then that same program reframed in positive terms). Even though we found clear support for the possibility that the negativity bias in reframing is object-specific, we cannot rule out the possibility that there also exists an object-general effect. We found tentative evidence for this possibility in Study 2, where we observed a small-to-medium negativity bias when participants evaluated a different program after reframing, $d = 0.38$, 95% CI [0.21, 0.55]. We observed a similar estimate of this effect in Study 1, $d = 0.23$, 95% CI [-0.12, 0.57]. Although this effect did not reach significance in Study 1, the larger sample size in Study 2 gave us greater power and precision to estimate the negativity bias in reframing in the different-program condition. Together, then, these two studies provide strong and consistent evidence for an object-specific effect, and tentative evidence for an object-general effect.

Given these results, we can speculate that future research and theorizing on the mechanisms underlying the negativity bias in reframing would do well to consider object-specific mechanisms as key drivers of these effects, while also allowing for the possibility that object-general processes play a (probably smaller) role. For example, past work on reframing has conceptualized the negativity bias as object-specific: Once an issue or person or event is described in negative terms, people construe it in negative terms and that negative construal tends to stick in the mind and resist subsequent attempts to change it (Ledgerwood & Boydston, 2014). Ledgerwood and Boydston (2014) drew an analogy to functional fixedness (Adamson, 1952; Duncker, 1945; Higgins & Chaires, 1980), suggesting that negative (vs. positive) mental labels may be especially likely to stick and resist reconceptualization. In a classic demonstration of functional fixedness, participants who were asked to solve a problem in which a box was used as a container later struggled to solve a task that required using the box as a shelf to support a candle

(Duncker, 1945). Notably, the second task was harder for these participants than for participants who had not initially conceptualized the box as a container, suggesting that the conceptualization of the box as a container inhibits the ability to reconceptualize it with a different function. It is plausible that frames—and especially negative frames—operate similarly. Given the present results and the centrality of subjective construal to social cognitive theory and research (Asch, 1940; Henderson & Wakslak, 2010; Lord & Lepper, 1999; Sánchez et al., 2021), we believe construal remains a promising contender for unpacking the mechanisms underlying sequential framing effects. Alternatively, it is plausible that sequential framing effects operate at the level of judgment, rather than representation.⁹ For example, perhaps people's preference for cognitive consistency (see e.g., Gawronski, 2012; Guadagno & Cialdini, 2010) leads them to avoid revising an initial judgment, especially when that judgment is negative. Future research could test this possibility.

In addition to illuminating the conditions that give rise to a negativity bias in reframing effects, the current research also provided the first clear evidence that both negative and positive frames are sticky, albeit to different extent. Past research has demonstrated that negative frames are relatively sticky compared to positive frames, but it was unclear whether positive frames can stick in the mind, too. In our experiments, initial valenced frames—both negative and positive—resulted in less change in attitude toward a policy program when it was subsequently reframed, compared to attitude toward a different policy program. These results are consistent with the possibility that an initial valenced frame may lead people to conceptualize an object as fundamentally bad or good, and that conceptualization may stick in the mind and resist subsequent reconceptualization. This resistance against reconceptualization is particularly strong for objects that are framed negatively first: The stain of negativity is particularly difficult to counteract. Yet an initial positive conceptualization can still buttress against the effect of a subsequent negative reframe to some extent.

LIMITATIONS AND FUTURE DIRECTIONS

Our results suggest that the negativity bias in reframing documented in past research (Boydston et al., 2019; Itani et al., 2023; Ledgerwood & Boydston, 2014; Sparks & Ledgerwood, 2019) operates primarily in an object-specific manner. Thus, efforts to mitigate this negativity bias may be more successful when they focus on promoting reconceptualization of the target stimulus as opposed to promoting a more positive general outlook. As in past research on reframing, however, the frames we used in the current research were relatively mundane and unlikely to evoke strong affective responses. Given research suggesting that negative affect is stronger and longer lasting than positive affect (Brosschot & Thayer, 2003; Costa et al., 2014; Scherer & Wallbott, 1994), it seems plausible that more affectively laden frames, such as those that evoke fear, sadness, or anger, could have a more diffuse, object-general effect compared to the ones used in the present research.

9. We thank Hans Alves for suggesting this possibility.

In Study 2, we examined whether rebranding an object by giving it a different name was sufficient to reduce the negativity bias in reframing, and we found that it was not. However, we found tentative evidence that rebranding may promote attitude change toward Time 2 frame regardless of frame valence order. This finding suggests that rebranding may be equally effective at enabling positive and negative reconceptualization and get people “unstuck” from the influence of an initial valenced frame. What’s in a name, then? Not much when it comes to mitigating negativity bias, it seems, but perhaps a more general potential to enable greater attitude change toward a new frame. The robustness of this finding awaits future confirmatory research. Future research can also zero in on the reason why rebranding seems to allow for greater attitude change across valenced frames (e.g., rebranding might afford inferences that the essence of an object has changed, allowing people to see it in a new light).

As recommended by Simons et al. (2017) and the Society for Personality and Social Psychology (2023), we highlight limitations that constrain the generality of the present findings. We expect that our findings will generalize to similar samples of U.S. adults. However, we think it is especially important to delineate our sample in terms of age, culture, and race, to avoid overclaiming beyond the borders of our sample population. First, there may be age-related constraints on the present findings. Though the sample contained a relatively wide age range (S1: $M_{\text{age}} = 38.3$, $SD_{\text{age}} = 12.0$; S2: $M_{\text{age}} = 37.7$, $SD_{\text{age}} = 12.6$), Sparks and Ledgerwood (2019) found that the size of the negativity bias in reframing declines as age increases. Therefore, future research should investigate whether the effects that we document here also change with increasing age, a possibility that could explain the age-related attenuation of the negativity bias.

In addition, research on cross-cultural differences in social cognition indicates that we should not assume the present findings will generalize to other cultural contexts (Lechuga et al., 2011; Nisbett et al., 2001; San Martin et al., 2018), nor should we assume that findings observed in a primarily White sample will necessarily generalize to participants whose life experiences have not been shaped by societal hierarchies in the same way (Remedios et al., 2022; Roberts et al., 2020). As our sample is U.S.-specific and primarily White, we should not assume that our results will generalize beyond these boundaries. Instead, further research is needed to explore potential mechanisms in additional cultural contexts and beyond primarily White samples (see Ledgerwood et al., in press).

Finally, future studies should assess whether our characterization of reframing effects as object-specific extends to reframing effects in the gain domain (i.e., when people contemplate the presence or absence of rewards rather than the presence or absence of costs). As reframing effects appear to function differently in the loss versus gain domain, we do not expect that the pattern of results identified here will necessarily be the same in the gain domain (Sparks & Ledgerwood, 2017; see also Sun et al., 2022). However, we do anticipate that the findings will generalize to other loss-domain scenarios, as reframing research shows consistent results across a range of scenarios (Sparks & Ledgerwood, 2019).

CONCLUSION

Taken together, the current research provides the first direct evidence that the negativity bias in reframing observed in past research is largely object-specific. Furthermore, it suggests for the first time that positive frames are sticky, too—just less so than negative frames. Thus, at least in the samples and contexts studied here, both positive and negative frames tend to stick in the mind, but negative frames are especially sticky; they follow the particular objects to which they are first applied, and their greater stickiness cannot be erased via superficial rebranding.

Author contributions. Y. Andre Wang: methodology, formal analysis, validation, resources, data curation, writing—original draft and review and editing, visualization, supervision, project administration; Melissa C. Liwag: methodology, investigation, formal analysis, writing—original draft; Katherine Weltzien: writing—original draft, supervision, project administration; Trevor Crowell: writing—original draft and review and editing; Alison Ledgerwood: conceptualization, methodology, resources, supervision, writing—original draft and review and editing.

Acknowledgments. This research was partially supported by a Social Sciences and Humanities Research Council (SSHRC) Insight Development Grant (430-2022-00087) to Y. Andre Wang.

Supplemental materials. Materials, datasets, and data analysis scripts for all studies are available at <https://osf.io/6gec8/> (Wang et al., 2022).

REFERENCES

- Adamson, R. E. (1952). Functional fixedness as related to problem solving: A repetition of three experiments. *Journal of Experimental Psychology*, *44*(4), 288–291.
- Asch, S. E. (1940). Studies in the principles of judgments and attitudes: II. Determination of judgments by group and by ego standards. *Journal of Social Psychology*, *12*(2), 433–465. <https://doi.org/10.1080/00224545.1940.9921487>
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*(3), 258–290. <https://doi.org/10.1037/h0055756>
- Boydston, A. E., Ledgerwood, A., & Sparks, J. (2019). A negativity bias in reframing shapes political preferences even in partisan contexts. *Social Psychological and Personality Science*, *10*(1), 53–61. <https://doi.org/10.1177/1948550617733520>
- Brosschot, J. F., & Thayer, J. F. (2003). Heart rate response is longer after negative emotions than after positive emotions. *International Journal of Psychophysiology*, *50*(3), 181–187. [https://doi.org/10.1016/s0167-8760\(03\)00146-6](https://doi.org/10.1016/s0167-8760(03)00146-6)
- Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin*, *25*(2), 215–224. <https://doi.org/10.1177/0146167299025002007>
- Costa, T., Cauda, F., Crini, M., Tatu, M. K., Celegghin, A., de Gelder, B., & Tamietto, M. (2014). Temporal and spatial

- neural dynamics in the perception of basic emotions from complex scenes. *Social Cognitive and Affective Neuroscience*, 9(11), 1690–1703. <https://doi.org/10.1093/scan/nst164>
- Cramer, A. O., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P., Waldorp, L. J., & Wagenmakers, E. J. (2006). Hidden multiplicity in exploratory multiway ANOVA: Prevalence and remedies. *Psychonomic Bulletin and Review*, 23(2), 640–647. <https://doi.org/10.3758/s13423-015-0913-5>
- da Silva Frost, A., & Ledgerwood, A. (2020). Calibrate your confidence in research findings: A tutorial on improving research methods and practices. *Journal of Pacific Rim Psychology*, 14, e14. <https://doi.org/10.1017/prp.2020.7>
- da Silva Frost, A., Wang, Y. A., Eastwick, P. W., & Ledgerwood, A. (2022). Summarized attribute preferences have unique antecedents and consequences. *Journal of Experimental Psychology: General*. Advance online publication. <https://doi.org/10.1037/xge0001242>
- Duncker, K. (1945). On problem-solving (L. S. Lees, Trans.). *Psychological Monographs*, 58(5, Whole No. 270).
- Fang, X., van Kleef, G. A., & Sauter, D. A. (2018). Person perception from changing emotional expressions: Primacy, recency, or averaging effect? *Cognition & Emotion*, 32(8), 1597–1610. <https://doi.org/10.1080/02699931.2018.1432476>
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/bf03193146>
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50(2), 229–238. <https://doi.org/10.1037//0022-3514.50.2.229>
- Frick, R. W. (1998). A better stopping rule for conventional statistical tests. *Behavior Research Methods, Instruments, & Computers*, 30, 690–697.
- Gawronski, B. (2012). Back to the future of dissonance theory: Cognitive consistency as a core motive. *Social Cognition*, 30(6), 652–668. <https://doi.org/10.1521/soco.2012.30.6.652>
- Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General*, 139(4), 683–701. <https://doi.org/10.1037/a0020315>
- Guadagno, R. E., & Cialdini, R. B. (2010). Preference for consistency and social influence: A review of current research findings. *Social Influence*, 5(3), 152–163. <https://doi.org/10.1080/15534510903332378>
- Hall, E. V., Phillips, K. W., & Townsend, S. S. (2015). A rose by any other name?: The consequences of subtyping “African-Americans” from “Blacks”. *Journal of Experimental Social Psychology*, 56, 183–190. <https://doi.org/10.1016/j.jesp.2014.10.004>
- Henderson, M. D., & Waksalak, C. J. (2010). Over the hills and far away: The link between physical distance and abstraction. *Current Directions in Psychological Science*, 19(6), 390–394. <https://doi.org/10.1177/0963721410390802>
- Higgins, E. T., & Chaires, W. M. (1980). Accessibility of interrelational constructs: Implications for stimulus encoding and creativity. *Journal of Experimental Social Psychology*, 16(4), 348–361. [https://doi.org/10.1016/0022-1031\(80\)90027-X](https://doi.org/10.1016/0022-1031(80)90027-X)
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1), 1–55. [https://doi.org/10.1016/0010-0285\(92\)90002-J](https://doi.org/10.1016/0010-0285(92)90002-J)
- Isaac, M. (2021, October 28). Facebook renames itself Meta. *New York Times*. <https://www.nytimes.com/2021/10/28/technology/facebook-meta-name-change.html>
- Itani, A. I., Coleman, C. W., AlGhazali, R., AlMalik, M., da Silva Frost, A., Fedavi, N., Imran, M., Weltzien, K., Yousef, S., Ledgerwood, A., & Maitner, A. T. (2023). Are negative frames equally sticky across cultural contexts? Exploring sequential framing effects with Arab participants in the UAE. *Current Research in Ecological and Social Psychology*, 5, 100129. <https://doi.org/10.1016/j.cresp.2023.100129>
- Ito, T. A., Larsen, J. T., Smith, N. K., & Cacioppo, J. T. (1998). Negative information

- weighs more heavily on the brain: The negativity bias in evaluative categorizations. *Journal of Personality and Social Psychology*, 75(4), 887–900. <https://doi.org/10.1037//0022-3514.75.4.887>
- Klein, N., & O'Brien, E. (2016). The tipping point of moral change: When do good and bad acts make good and bad actors? *Social Cognition*, 34(2), 149–166. <https://doi.org/10.1521/soco.2016.34.2.149>
- Krishnamurthy, P., Carter, P., & Blair, E. (2001). Attribute framing and goal framing effects in health decisions. *Organizational Behavior and Human Decision Processes*, 85(2), 382–399. <https://doi.org/10.1006/obhd.2001.2962>
- Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, 44(7), 701–710. <https://doi.org/10.1002/ejsp.2023>
- Lechuga, J., Swain, G. R., & Weinhardt, L. S. (2011). The cross-cultural variation of predictors of human papillomavirus vaccination intentions. *Journal of Women's Health*, 20(2), 225–230.
- Ledgerwood, A., & Boydston, A. E. (2014). Sticky prospects: Loss frames are cognitively stickier than gain frames. *Journal of Experimental Psychology: General*, 143(1), 376–385. <https://doi.org/10.1037/a0032310>
- Ledgerwood, A., da Silva Frost, A., Kadirvel, S., Maitner, A., Wang, Y. A., & Maddox, K. B. (in press). Methods for advancing an open, replicable, and inclusive science of social cognition. In K. Hugenberg, K. Johnson, & D. E. Carlston (Eds.), *The Oxford handbook of social cognition*.
- Ledgerwood, A., Mandisodza, A. N., Jost, J. T., & Pohl, M. J. (2011). Working for the system: Motivated defense of meritocratic beliefs. *Social Cognition*, 29(3), 322–340. <https://doi.org/10.1521/soco.2011.29.3.322>
- Ledgerwood, A., Soderberg, C. K., & Sparks, J. (2017). Designing a study to maximize informational value. In J. Plucker & M. Makel (Eds.), *Toward a more perfect psychology: Improving trust, accuracy, and transparency in research* (pp. 33–58). American Psychological Association.
- Levin, I. P., & Gaeth, G. J. (1988). How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research*, 15(3), 374–378. <https://doi.org/10.1086/209174>
- Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, 76(2), 149–188. <https://doi.org/10.1006/obhd.1998.2804>
- Lord, C. G., & Lepper, M. (1999). Attitude representation theory. *Advances in Experimental Social Psychology*, 31, 265–343. [https://doi.org/10.1016/S0065-2601\(08\)60275-0](https://doi.org/10.1016/S0065-2601(08)60275-0)
- Mao, A., Barnes, K., Sharpe, L., Geers, A. L., Helfer, S. G., Faasse, K., & Colagiuri, B. (2021). Using positive attribute framing to attenuate nocebo side effects: A cybersickness study. *Annals of Behavioral Medicine*, 55(8), 769–778. <https://doi.org/10.1093/abm/kaaa115>
- Markus, H., & Zajonc, R. B. (1985). The cognitive perspective in social psychology. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology* (3rd ed., pp. 137–230). Random House.
- McShane, B. B., Gal, D., Gelman, A., Robert, C., & Tackett, J. L. (2019). Abandon statistical significance. *American Statistician*, 73(Suppl 1), 235–245. <https://doi.org/10.1080/00031305.2018.1527253>
- Muzellec, L., & Lambkin, M. (2006). Corporate rebranding: Destroying, transferring or creating brand equity? *European Journal of Marketing*, 40(7/8), 803–824. <https://doi.org/10.1108/03090560610670007>
- Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic versus analytic cognition. *Psychological Review*, 108(2), 291–310. <https://doi.org/10.1037/0033-295x.108.2.291>
- Ohira, H., Winton, W. M., & Oyama, M. (1998). Effects of stimulus valence on recognition memory and endogenous eyeblinks: Further evidence for positive-negative asymmetry. *Personality and Social Psychology Bulletin*, 24(9), 986–993. <https://doi.org/10.1177/0146167298249006>
- Pahl, R. (2022). GroupSeq: Group sequential design probabilities—with graphical user interface (Version 1.4.2) [Computer

- software]. <https://CRAN.R-project.org/package=GroupSeq>
- Paul, K. (2021, October 28). Facebook announces name change to Meta in rebranding effort. *The Guardian*. <https://www.theguardian.com/technology/2021/oct/28/facebook-name-change-rebrand-meta>
- Payne, K., & Lundberg, K. (2014). The affect misattribution procedure: Ten years of evidence on reliability, validity, and mechanisms. *Social and Personality Psychology Compass*, 8(12), 672–686. <https://doi.org/10.1111/spc3.12148>
- Philogène, G. (2001). Stereotype fissure: Katz and Braly revisited. *Social Science Information*, 40(3), 411–432. <https://doi.org/10.1177/053901801040003003>
- Pietri, E. S., Fazio, R. H., & Shook, N. J. (2013a). Recalibrating positive and negative weighting tendencies in attitude generalization. *Journal of Experimental Social Psychology*, 49(6), 1100–1113. <https://doi.org/10.1016/j.jesp.2013.08.001>
- Pietri, E. S., Fazio, R. H., & Shook, N. J. (2013b). Weighting positive versus negative: The fundamental nature of valence asymmetry. *Journal of Personality*, 81(2), 196–208. <https://doi.org/10.1111/j.1467-6494.2012.00800.x>
- Quattrone, G. A., & Tversky, A. (2004). Contrasting rational and psychological analyses of political choice. In E. Shafir (Ed.), *Preference, belief, and similarity: Selected writings by Amos Tversky* (pp. 621–644). Boston: Review.
- Remedios, J. D. (2022). Psychology must grapple with Whiteness. *Nature Reviews Psychology*, 1, 125–126. <https://doi.org/10.1038/s44159-022-00024-4>
- Roberts, S. O., Bareket-Shavit, C., Dollins, F. A., Goldie, P. D., & Mortenson, E. (2020). Racial inequality in psychological research: Trends of the past and recommendations for the future. *Perspectives on Psychological Science*, 15(6), 1295–1309. <https://doi.org/10.1177/1745691620927709>
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5(4), 296–320. https://doi.org/10.1207/S15327957PSPR0504_2
- Sagarin, B. J., Ambler, J. K., & Lee, E. M. (2014). An ethical approach to peering at data. *Perspectives on Psychological Science*, 9(3), 293–304. <https://doi.org/10.1177/1745691614528214>
- San Martin, A., Sinaceur, M., Madi, A., Tompson, S., Maddux, W. W., & Kitayama, S. (2018). Self-assertive interdependence in Arab culture. *Nature Human Behaviour*, 2, 830–837. <https://doi.org/10.1038/s41562-018-0435-z>
- Sánchez, A. M., Coleman, C. W., & Ledgerwood, A. (2021). Does temporal distance influence abstraction? A large pre-registered experiment. *Social Cognition*, 39(3), 352–365. <https://doi.org/10.1521/soco.2021.39.3.352>
- Scherer, K. R., & Wallbott, H. G. (1994). Evidence for universality and cultural variation of differential emotion response patterning. *Journal of Personality and Social Psychology*, 66(2), 310–328. <https://doi.org/10.1037//0022-3514.66.2.310>
- Sia, T. L., Lord, C. G., Blessum, K. A., Ratcliff, C. D., & Lepper, M. R. (1997). Is a rose always a rose? The role of social category exemplar change in attitude stability and attitude-behavior consistency. *Journal of Personality and Social Psychology*, 72(3), 501–514. <https://doi.org/10.1037/0022-3514.72.3.501>
- Simons, D. J., Shoda, Y., & Lindsay, D. S. (2017). Constraints on Generality (COG): A proposed addition to all empirical papers. *Perspectives on Psychological Science*, 12(6), 1123–1128. <https://doi.org/10.1177/1745691617708630>
- Smith, E. A., & Malone, R. E. (2003). Altria means tobacco: Philip Morris's identity crisis. *American Journal of Public Health*, 93(4), 553–556. <https://doi.org/10.2105/ajph.93.4.553>
- Society for Personality and Social Psychology. (2023). *Guidelines and examples for writing an effective constraints on generality section*. <https://spsp.org/professional-development/publishing-resources/resources-for-inclusive-practices/guidelines-and-examples-for-writing-an-effective-constraints-on-generality-section>
- Sørensen, K., Okan, O., Kondilis, B., & Levin-Zamir, D. (2021). Rebranding social distancing to physical distancing: Calling for a change in the health promotion vocabulary to enhance clear communication during a pandemic. *Global Health*

- Promotion*, 28(1):5–14. <https://doi.org/10.1177/1757975920986126>
- Sparks, J., & Ledgerwood, A. (2017). When good is stickier than bad: Understanding gain/loss asymmetries in sequential framing effects. *Journal of Experimental Psychology: General*, 146(8), 1086–1105. <https://doi.org/10.1037/xge0000311>
- Sparks, J., & Ledgerwood, A. (2019). Age attenuates the negativity bias in reframing effects. *Personality and Social Psychology Bulletin*, 45(7), 1042–1056. <https://doi.org/10.1177/0146167218804526>
- Sun, S., Hu, J., & Yu, R. (2022). Domain-specific neural substrates underlie the framing effect. *Neuroimage: Reports*, 2(4), 100119. <https://doi.org/10.1016/j.ynrp.2022.100119>
- Wang, Y. A., Liwag, M., Weltzien, K., Crowell, T., & Ledgerwood, A. (2022, September 9). Sticky frames and what's in a name: Frames stick to particular objects. <https://osf.io/6gec8/>
- Whitmarsh, L. (2009). What's in a name? Commonalities and differences in public understanding of "climate change" and "global warming." *Public Understanding of Science*, 18(4), 401–420. <https://doi.org/10.1177/0963662506073088>