Worse is Bad: Divergent Inferences From Logically Equivalent Comparisons

Yoel Inbar
University of Toronto

Ellen Evers
Haas School of Business, University of California, Berkeley

Article in press at *Journal of Experimental Psychology: General*.

Address correspondence to Yoel Inbar, Department of Psychology, University of Toronto, yoel.inbar@utoronto.ca. Data, materials, and Stata code to reproduce the reported analyses are available at <https://osf.io/5kn2r/>.

Abstract

Comparisons can often be framed in different but equivalent ways. For example, “A is better than B,” can also be expressed as “B is worse than A.” In seven studies (and four further studies in the Supplemental Materials) we find that logically-equivalent comparison frames have divergent effects on judgments and choices for the items being compared as well as other members of the set from which those items were drawn. These effects are asymmetric, affecting inferior items more strongly than superior ones. We propose a “comparison framing” account that draws on theory in linguistics on the “markedness” of adjectives (Cruse, 1976; Lehrer, 1985) to explain these results. We show that this account fits the data better than two previous accounts of attribute framing: automatic valenced associations (Levin, Schneider, & Gaeth, 1998) and leakage of information (McKenzie & Nelson, 2003).

Worse is Bad: Divergent Inferences from Logically Equivalent Comparisons

Imagine that there is a new restaurant in town that you are interested in trying. You ask your food-enthusiast friend whether she's visited. She tells you, "I've been there twice, and it was mixed—the service was always good, but one time, my entrée was better than the other." How interested would you be in trying the restaurant?

Now imagine that your friend had instead told you, "one time, my entrée was worse than the other." Would this change your answer? These two statements are, of course, logically equivalent—to say that the entrée was better on one visit logically implies that it was worse on the other. But people may not treat them equivalently. If you are like the authors, the options being compared—and, particularly, the inferior option—sound more appealing under the former description ("better") than the latter ("worse"). These differences in evaluations may even extend beyond the two items being compared to the set from which they were drawn. In this case, it may be that the impression of other dishes available at that restaurant is less positive following the "worse" comparison.

Linguistic markedness and comparison frames

We propose that research from linguistics on antonym "markedness" (Bolinger, 1977; Clark, 1969a; Cruse, 1966; Greenberg, 1966; Lehrer, 1985, 2002) can explain how logically-equivalent comparisons might lead to different judgments. According to this research, *gradable antonyms* (i.e., antonyms that refer to opposite sections of an underlying value scale, such as good/bad, tall/short, or hot/cold) can be either "marked" or "unmarked." For our purposes, the most important distinction between them is that only marked antonyms carry a directional presumption. For example, consider the antonym pair "long" (unmarked) and "short" (marked). Asking "how long is the play" does not presume that it is long, whereas "how short is the play"

presumes that it is short. Because they do not presume a scale position, unmarked members of antonym pairs can be used in more contexts. For example, “90 minutes short” sounds more unusual to English speakers than “90 minutes long.” According to researchers in linguistics (e.g., Lehrer, 2002), all antonym pairs have at least one marked member. For most, one member is marked and the other is unmarked; for some, both members are marked (e.g., hot/cold).

Markedness is a descriptive theory of how language is used, not of its psychological effects. However, psychological researchers have found that markedness also has consequences for people’s judgments. For example, asking participants to make estimates using unmarked or marked antonyms—e.g., “How high was the plane flying” (unmarked) vs. “How low was the plane flying?” (marked)—showed higher variance in estimates with the unmarked antonym, consistent with it not carrying a directional presumption (Harris, 1973). The theory developed by Clark (1969a) proposes that because unmarked members of an antonym pair can be used to refer to a scale of measurement without implying a scale position (Clark calls this the “nominal” sense of the antonym), they are simpler to process, store, and retrieve from memory. Marked antonyms can only be used in a “contrastive” sense that also encodes a scale position—that is, they contain an extra element compared to their unmarked counterparts. (This is why “How short is the play” is not a neutral question—“short” always encodes a position on the length scale.) There is substantial empirical evidence for a processing advantage for unmarked antonyms. Clark (1969a, 1969b) has found that people process unmarked antonyms more quickly, remember relations encoded with them better, and solve three-term comparison problems stated with unmarked antonyms more quickly and accurately (e.g., “John is taller than Sue, and Sue is taller than Jim. Is John taller than Jim?”). Likewise, when evaluating whether images and verbal descriptions match, marked comparatives (B is below A) are processed more slowly than unmarked

comparatives (A is above B) (Chase & Clark, 1972; Clark & Chase, 1972; see also Trabasso, Rollins, & Shaughnessy, 1971).

We propose that in addition to its other effects on judgment, markedness has implications for the interpretation of evaluative comparisons (such as “the steak was better than the pasta”). “Better” and “worse” are the comparative forms of “good” (unmarked) and “bad” (marked) (Cruse, 1976; Lehrer, 1985). Because “worse” is marked, it encodes both a scale (quality) and a position (low). Therefore, the comparison “B is worse than A” should lead people to see both A and B as lower-quality than the logically-equivalent comparison “A is better than B.” However, marked comparisons have an extra feature that we predict will also affect judgments. Theorists in linguistics (e.g., Bollinger, 1977; Cruse, 1976; Lehrer, 1985, 2002) have argued that for many pairs of evaluative comparatives—including “better” and “worse”—antonym pairs are *asymmetrically committed*. According to this theory, an entity described using the marked member is “committed” to the negative end of the scale. In other words, the comparison “B is worse than A” entails that B *must* be bad, and that A *may* be as well (Lehrer, 1985, 2002). This predicts asymmetric effects of markedness on judgment for evaluative comparisons, such that effects will be stronger for the item described as inferior (i.e., the item “committed” by the marked comparative) than for the item described as superior (i.e., the uncommitted item).

Asymmetric commitment also predicts effects of markedness for items beyond those being compared directly. Because a “worse” comparison entails that one of the items being compared is bad, whereas a “better” comparison does not, negatively-framed comparisons should extend the bottom of the perceived quality range (i.e., “this set contains some bad items”). We predict that this will shift perceptions of other set members at the bottom of the quality range

downwards. That is, we expect that negatively-framed evaluative comparisons will lead to lower quality judgments of other items at the bottom of the quality range.

This account, which we call “comparison framing,” predicts that logically-equivalent evaluative comparisons will have divergent effects on judgment and choice, and makes specific predictions about the form of these differences. We next review how extant research on framing contrasts with the account we have proposed.

Previous Accounts of Framing Effects

The distinction we have described between “better” and “worse” comparisons is a specific instance of a more general fact: A single attribute (or relationship between attributes) can often be described in different but logically equivalent ways (where logically equivalent means simply that one description necessarily entails the other; Sher & McKenzie, 2006). For example, if someone likes chocolate ice cream more than strawberry ice cream, by definition they like strawberry ice cream less than chocolate. If 95% of patients survive a medical procedure, by definition 5% die. If a policy allocates more of a finite resource to people in group A, by definition it allocates less of that resource to people not in group A.

Framing effects occur when two logically equivalent descriptions reliably lead to different judgments or choices, violating normative models of rational choice (Tversky & Kahneman, 1981). This is a description of an outcome, not an explanation of the psychology underlying it, and there have been different psychological accounts proposed to account for framing effects. Early studies investigating framing did so for decisions under uncertainty (decisions involving probabilistic outcomes; see Tversky & Kahneman, 1981). Later, framing effects were observed for non-probabilistic outcomes. The most well-known type of non-probabilistic framing has been called “attribute framing” (for a review see Levin, Schneider, &

Gaeth, 1996). In attribute framing studies, targets are described on an attribute using positive or negative language, which affects subsequent judgments. For example, when people read about ground beef that is either described as “75% lean” or “25% fat,” they expect that it will be better-quality, less greasy, and better-tasting under the “75% lean” frame. These framing effects are thought to result from automatic valenced associations evoked by the frames. Under this account, positive framing “leads to an encoding of the information that tends to evoke favorable associations in memory,” whereas negative framing “is likely to cause an encoding that evokes unfavorable associations” (Levin, Schneider, & Gaeth, 1998, p. 164). The positive or negative encoding of the information is then thought to lead to biased attention to and retrieval of valence-congruent information about the framed item (see also Levin, 1987; Levin & Gaeth, 1988).

The first difficulty in applying this research to comparisons is that it concerns evaluations of single items, and comparisons by definition involve multiple items. However, following the logic of the single-item studies, perhaps positive or negative comparisons evoke valenced associations that cause attention to and retrieval of valence-congruent information about both items. This would predict that both items would be evaluated more positively when described with an unmarked comparative (“better”) vs. a marked comparative (“worse”). However, in contrast to the comparison framing account, the valenced associations account would not predict stronger effects for the inferior item. It might predict equal effects for other items in the choice set consistent with the valence of the associations, but not specific effects for some other items in the set. Therefore, the valenced associations account makes different predictions from the comparison framing account.

Leakage of information

An alternative interpretation of attribute framing effects has been proposed by McKenzie and colleagues (Leong, McKenzie, Sher, & Müller-Trede, 2017; McKenzie & Nelson, 2003; Sher & McKenzie, 2006, 2008), who argue that they may result from the normatively-appropriate use of information leaked by the communicator's choice of frame. According to the leakage of information account, a communicator who describes ground beef as 25% fat (compared to 75% lean) may have *chosen* to do so to signal that the meat is more fat than normal. More precisely, the account holds that “the frame [...] chosen by speakers to describe a current situation tends to be the one that has increased relative to their reference point” (McKenzie & Nelson, 2003, p. 597). For example, “ground beef would be more likely to be described in terms of 25% *fat* if most ground beef had less, rather than more, fat” (McKenzie & Nelson, 2003, p. 597, emphasis in original). If communicators' choices among logically equivalent frames can be informative about their beliefs, the frames may not be informationally equivalent, and thus treating them non-equivalently can be normatively defensible.

This account is intended to explain results from attribute framing studies in which a target is described in terms of a specific numeric quantity (e.g., 25% fat, 50% success rate, 5% mortality rate). It does not readily extend to comparison frames, where there is no quantity to exceed a reference point. When one brand of ground beef is described as worse than another, what would the reference point be, and in what sense could the “worse” frame include a quantity that has “increased” relative to it? As specified, the psychological mechanism posited by the leakage of information account does not appear to directly apply to comparison framing. One could modify the mechanism to apply to both items being compared (e.g., “both the items being compared are below the average quality”). In this case, the prediction would be that comparison frames should affect evaluations of the superior and inferior items equally. It would not predict

asymmetric inferences, nor would it predict effects for any other items in the choice set.

Therefore, the leakage of information account does not make the predictions that the comparison framing account does.

These predictions are derived from the specific mechanisms posited by the leakage of information account (e.g., McKenzie & Nelson, 2003). However, one can also construe the idea of leakage of information much more broadly, to hold that people go beyond what is stated explicitly to infer information from a communicator's utterances. The comparison framing account is compatible with a version of this idea, although it more precisely specifies which kind of information is inferred from comparisons, and under which circumstances. We return to the similarities and differences between the two accounts in the General Discussion.

The current studies

We examine the effects of comparison frames in seven studies. Studies 1a and 1b demonstrate the effect of comparison framing and its consequences for preferences over the compared items. Study 2 replicates these results and investigates to what extent comparison frames affect judgments of other items in the choice set, a context in which our theory predicts effects for other items, but extant accounts of framing do not. Study 3 tests a condition under which our account predicts that the effect of comparison frame should be equally strong for both items. Study 4 makes the communicator's choice of frame uninformative and finds that the effect of comparison framing on judgment persists even when it is clear that no information should be derived from the choice of comparison frame. Studies 5 and 6 test the consequences of comparison framing on choice. In Study 5 the choices are hypothetical; in Study 6, they are consequential (i.e., incentivized).

In all studies, we report all measures, conditions, and data exclusions. Sample sizes were always determined in advance based on cost and (when available) effect sizes from other studies. All data, code, and materials are available at <https://osf.io/5kn2r>.

Study 1a

In Study 1a, we tested the effect of comparison framing on evaluations of the two items being compared. Participants read about two toasters, both of which had been selected by Consumer Reports as in the top 10 of all toasters tested. In “better” condition, participants read that one of the two (the KitchenAid) had been described as “significantly better” than the other (the JC Penney). In the “worse” condition, participants read that the JC Penney toaster had been described as “significantly worse” than the KitchenAid. The descriptions were otherwise identical. Participants were then asked to rate each toaster (in counterbalanced order) on 11-point scales (-5 = “extremely bad”; 0 = “middle of the pack”; 5 = “extremely good”). After the rating questions, participants were asked to indicate whether they would prefer to receive the JC Penney toaster or another toaster randomly selected from the top 25 toasters rated by Consumer Reports. Finally, participants answered a memory check question that asked them to recall which toaster Consumer Reports preferred (this question included a “Don’t remember” option).

The sample size and analysis plan for this study were preregistered at <https://aspredicted.org/blind.php?x=3x9tz7>. All analytic choices in this study can be considered confirmatory unless otherwise specified. We set a target sample size of 800 participants from Prolific Academic and obtained 801 complete responses. Of these, 609 (304 in the “worse” condition and 305 in the “better” condition) correctly answered the memory check question. In this and the following studies, we report analyses including only participants who pass memory

checks, but unless otherwise noted, the direction and significance of results do not change if all participants are included.

Participants rated the inferior (JC Penney) toaster significantly more negatively in the “worse” condition as compared to the “better” condition, $M_{better} = .92$, $SD_{better} = 1.93$; $M_{worse} = -1.17$, $SD_{worse} = 2.30$; $t(607) = 12.12$, $p < .001$, $d = .98$. This was also for the case for the superior (KitchenAid) toaster, albeit less strongly, $M_{better} = 3.49$, $SD_{better} = 1.26$; $M_{worse} = 2.88$, $SD_{worse} = 1.48$; $t(607) = 5.47$, $p < .001$, $d = .44$ (see Figure 1, left panel). In an ANOVA with rating target as a within-subjects factor and condition as a between-subjects factor, the interaction between rating target and condition was significant, $F(1, 607) = 68.92$, $p < .001$, indicating a significantly stronger effect of condition on ratings of the inferior (JC Penney) toaster.

The choice data were consistent with the rating scales. Participants in the “worse” condition were more likely to prefer another toaster from the top 25 over the JC Penney toaster (206/304; 67.8%) than were participants in the “better” condition (133/305; 43.6%), Pearson $\chi^2(1) = 36.0$, $p < .001$.

The results of Study 1a are consistent with the comparison framing account but not existing accounts of framing effects. Both items were rated as significantly worse under the “worse” comparison frame compared to the “better” frame, but this effect was stronger for the inferior item (this asymmetry is not predicted by existing accounts). In Study 1b, we test both hypotheses again, using a different set of stimuli and a different evaluative comparison.

Figure 1. Quality ratings from Study 1a (left panel) and Study 1b (right panel) by comparison frame condition. In Study 1a, the Kitchen Aid toaster was described as superior to the JC Penney toaster; in Study 1b, Team A was described as superior to Team B. Ratings are on scales ranging from -5 to 5 (labelled as “extremely bad/extremely good” in Study 1a and “extremely unqualified/extremely qualified” in Study 1b). Error bars are 95% confidence intervals.

Study 1b

In Study 1b, we tested the generalizability of comparison framing in a new setting. We also tested whether the framing effects we observed depended on the use of the specific terms “better” and “worse,” or whether they would generalize to a different evaluative comparison. Participants were asked to imagine that they were a manager at a large technology firm overseeing an important new project. Their assistant had compiled a list of the top 10 engineering teams at the firm. Of these, only two could start soon enough to keep the project on schedule: Team A and Team B. The assistant described as rating Team A as “significantly more qualified” than Team B, or Team B as “significantly less qualified” than Team A. Participants were then asked to rate both teams (in counterbalanced order) on 11-point scales (-5 = “extremely unqualified”; 0 = “middle of the pack”; 5 = “extremely qualified”). After the rating questions, participants were asked to imagine that Team A was not available after all, and to choose whether they would prefer to go with Team B, or to choose another team from the top 50 at the company. (Note that as the scenario had stated that Teams A and B were the only ones available in the top 10, this meant that the replacement team would not be in the top 10.) Finally, participants answered two memory check questions asking them to indicate whether each team was in the assistant’s top 10 list (response options were “Yes,” “No,” and “Don’t Remember”).

We set a target sample size of 450 participants from Amazon Mechanical Turk and obtained 452 complete responses. We retained for analysis the 313 participants who correctly answered both memory check questions. Of these, 156 had been randomly assigned to the “less qualified” condition and 157 to the “more qualified” condition. Participants rated the inferior (B) team significantly more negatively in the “less qualified” condition as compared to the “more qualified” condition, $M_{more} = 1.27$, $SD_{more} = 1.67$; $M_{less} = -.026$, $SD_{less} = 2.47$; $t(311) = 5.46$, $p <$

.0001, $d = .62$. This was also the case for the superior (A) team, albeit less strongly, $M_{more} = 3.85$, $SD_{more} = .99$; $M_{less} = 3.39$, $SD_{less} = 1.45$; $t(311) = 3.29$, $p = .001$, $d = .37$ (see Figure 1, right panel). In an ANOVA with rating target as a within-subjects factor and condition as a between-subjects factor, the interaction between rating target and condition was significant, $F(1, 311) = 11.81$, $p < .001$, indicating a significantly stronger effect of condition on ratings of the inferior (B) team.

The choice data were consistent with the rating scales. Participants in the “less qualified” condition were more likely to prefer another team from the top 50 over Team B (68/156; 43.59%) than were participants in the “more qualified” condition (32/157; 20.38%), Pearson $\chi^2(1) = 19.38$, $p < .0001$.

Discussion

Study 1b replicated the results of Study 1a: Evaluations of both Team A and Team B were more negative when they were compared using the “less qualified” framing than the (logically equivalent) “more qualified” framing. Again, however, this effect was asymmetric; judgments of Team B (the inferior team) were more affected than judgments of Team A.

These results also show that the effects of comparison framing generalize beyond the comparatives “better” and “worse.” Research in linguistics has focused on markedness as a characteristic of specific antonym pairs, but these results suggest that its psychological implications extend more broadly. We return to this point in the General Discussion.

Studies 1a and 1b both asked participants to evaluate items on Likert-type scales, raising the possibility that comparison framing effects in these studies are the result of participants’ confusion about how to use response scales (see e. g., Moore, 2007). Responses on subjective rating scales can reflect a combination of absolute and relative evaluations (Burson & Klayman,

2005; Moore, 2007) and as such it may be that subjective response scales invite participants to conflate relative evaluations (is A better than B?) and absolute evaluations (how good are A and B?). This account would not predict an asymmetry in the effects of framing on A vs. B, nor would it predict the consequences of comparison framing for choices that we observed in Studies 1a and 1b. Nevertheless, to further exclude this alternative explanation, we ran a study (Study S1, reported in the Supplemental Material) using percentile ranks rather than subjective response scales and found that comparison framing also affects responses expressed in percentiles.

In addition to asymmetric effects on evaluations of the compared items, our account also predicts effects of comparison framing on other members of the choice set. We next test the hypothesis that the negative comparison frame will extend the perceived lower end of the quality range downwards.

Study 2

Study 2 examined evaluations of other items in the choice set beyond the two being compared directly. In contrast to other accounts of framing, the comparison framing account predicts that a negative comparison frame will extend the perceived lower end of the set's quality range downwards, and so lead to more negative evaluations of the lowest-quality members of the set.

The experimental paradigm in the current study was very similar to Study 1b, except that we described Teams A and B as drawn from “a list of all teams at the company” (rather than from the top 10). In addition to rating how qualified they thought Teams A and B were, participants also rated 1) how qualified the average team at the company was; 2) how qualified the best team at the company was; 3) how qualified the worst team at the company was. Unlike

Study 1b, participants were not asked to make a choice between Team B and another unknown team.

We set a target sample size of 400 participants from Amazon Mechanical Turk and obtained 412 complete responses. We retained data from the 364 participants who correctly recalled that Team A had been rated as more qualified than Team B by the assistant. Of these, 171 participants had been randomly assigned to the “less qualified” condition and 193 to the “more qualified” condition. Descriptive statistics for ratings in each condition are shown in Table 1.

Ratings of Team B were significantly lower under the “less qualified” frame, $t(362) = 6.70, p < .0001, d = .70$; this was also true of ratings of Team A, but with a smaller effect size, $t(362) = 3.78, p = .002, d = .40$. Ratings of the “average” and “best” teams were not significantly different between conditions, both $ps > .20$. Ratings of the “worst” team were, however, significantly lower under the “less qualified” frame than the “more qualified” frame, $t(362) = 3.24, p = .001, d = .34$. These results are shown in Figure 2. In a mixed-model ANOVA with rating target as a within-subjects factor and condition as a between-subjects factor, the interaction between rating target and condition was significant, $F(4, 1448) = 10.0, p < .001$, indicating significant differences in the effect of condition on ratings across targets. Likewise, a mixed-model ANOVA on only ratings of Teams A and B only showed a significant interaction, $F(1, 362) = 10.36, p = .001$, indicating a significantly stronger effect of framing on Team B than Team A.

Discussion

The current results replicate Studies 1a-1b: comparison framing reliably affected perceptions of both items being compared, with stronger effects for the inferior than the superior

item. The current results also show that the effects of comparison frames extend beyond the two items directly being compared to other members of the set from which those items were drawn. Namely, the worst team at the company was seen as significantly less qualified under the “less qualified” frame than under the “more qualified” frame. We replicated this effect in two further studies described fully in the supplemental material (Studies S2 and S3). Study S2 ($N = 2,207$) was an exact replication of Study 2 conducted by a group of researchers not part of this project, and Study S3 ($N = 340$) was a close replication, differing only in minor wording changes. The results from both replications were very similar to those of Study 2. There were significant effects of framing for Team B and the “worst team,” smaller effects for Team A, and no effects for the “best” or “average” team. In both cases, the interactions between rating target and framing were significant, $ps < .001$.

Altogether, these results support the comparison framing account we have proposed over existing accounts. Although current accounts of attribute framing effects can be extended to predict main effects of comparison frames, they would not predict the asymmetric effects for inferior vs. superior items. Nor would they predict effects of framing for other lower-quality items in the choice set.

Table 1. *Means and standard deviations of team quality ratings by condition for Study 2. Means in bold differ significantly by condition.*

Comparison Frame	Team A <i>M (SD)</i>	Team B <i>M (SD)</i>	Best Team <i>M (SD)</i>	Average Team <i>M (SD)</i>	Worst Team <i>M (SD)</i>
More Qualified	3.37 (1.21)	-0.14 (1.66)	3.95 (1.24)	0.63 (1.12)	-1.72 (1.91)
Less Qualified	2.84 (1.50)	-1.44 (2.04)	3.78 (1.44)	0.50 (1.11)	-2.40 (2.15)

Figure 2. Quality ratings by comparison frame condition of Team A (the superior team), Team B (the inferior team), and the best, average, worst teams at the company. Ratings are on scales ranging from -5 (“extremely unqualified”) to 5 (“extremely qualified”). Error bars are 95% confidence intervals.

Study 3

Thus far, we have consistently found stronger effects of comparison frames on inferior target. We hypothesize that this asymmetry occurs because the negatively-framed comparison “commits” the inferior item (but not the superior item) to the negative end of the scale (see Lehrer, 1985, 2002). By this logic, committing *both* items to the negative end of the scale (e.g., by providing other information that makes it clear they are low-quality) should eliminate the asymmetry and lead to equally strong effects of comparison frames for both items. We tested this prediction in the current study.

We asked participants to imagine that they were managers who needed to lay off one of two teams working under them. Both these teams were described as “performing especially poorly” (thus committing them both to the negative end of the performance scale). However, participants saw different framings of the comparisons between the teams. Team A was either described as performing “significantly better” than Team B, or Team B was described as performing “significantly worse” than Team A. Participants were asked to rate how qualified they thought each team was on 11-point scales (-5 = “Extremely bad”; 0 = “Middle of the pack”; 5 = “Extremely good”). Finally, participants were asked to indicate whether Team A or B had been described as better-performing.

We set a target sample size of 400 participants from Amazon Mechanical Turk and obtained 410 complete responses. We retained data from the 370 participants who correctly recalled that Team A had been described as better-performing. Of these, 183 had been randomly assigned to the worse condition and 187 to the better condition. Replicating previous studies, ratings of Team B were significantly lower under the “worse” frame, $t(368) = 4.99, p < .0001, d = .52$. Unlike the previous studies, the effect on ratings of Team A was equally strong, $t(368) =$

4.75, $p < .0001$, $d = .49$. Descriptive statistics for ratings in each condition are shown in Table 2.

Consistent with an equal effect of condition on both ratings, in a mixed-model ANOVA with rating target as a within-subjects factor and condition as a between-subjects factor the interaction between rating target and condition was non-significant, $F(1, 368) = .54$, $p = .46$.¹

Table 2. Means and standard deviations by condition, Study 3.

Comparison Frame	Team A <i>M (SD)</i>	Team B <i>M (SD)</i>
Better	.53 (2.45)	-2.60 (2.01)
Worse	-.61 (2.19)	-3.58 (1.76)

The results of Study 3 show that comparison framing exerts an equally-strong effect on both items when they are known to be low-quality (i.e., both are “committed” to the bottom end of the scale). Notably, neither the previous asymmetry in evaluations nor its elimination in this study are predicted by existing accounts of framing effects.

Study 4

Having established that the robustness of the effects predicted by the comparison framing account, we examine whether communicator intentions moderate these effects. The leakage of information account (Leong, McKenzie, Sher, & Müller-Trede, 2017; McKenzie & Nelson, 2003; Sher & McKenzie, 2006, 2008) proposes that listeners may correctly derive information from the communicator’s choice of words. If the comparison framing effects we have observed thus far are the result of explicit inferences resulting from the communicator’s word choices,

¹ We also asked an exploratory choice question after the main evaluation measures. Participants were asked to imagine that it was not possible to lay off Team B and were asked to indicate whether they would choose to lay off Team A or find a different team to lay off. Consistent with the effect of condition on rating scale evaluations of Team A, a significantly higher percentage of participants chose to lay off Team A in the “worse” condition (73%) than in the “better” condition (59%).

making the choice of language uninformative about intentions should eliminate or at least attenuate any framing effects. In contrast, because markedness is an inherent feature of language, the comparison framing account predicts that judgments will not be moderated by information about speaker intentions—that, in other words, comparison framing will affect judgments even when the speakers’ use of one comparison or another does not convey anything about their beliefs.

Participants were asked to imagine that they were reviewing consumer evaluations of two new sodas (code named “soda A” and “soda B”) being tested by a market research company. In the *free response* condition, participants read that:

The market research company has been experimenting with different ways of asking people to evaluate products. In this case, they asked consumers to write out their evaluations.

Participants then saw written feedback from one consumer reading either “I think soda A tasted better than soda B” or “I think soda B tasted worse than soda A.” In the *form* condition, participants read that:

The market research company has been experimenting with different ways of asking people to evaluate products. In this case, focus group participants were asked to fill out pre-printed forms that asked EITHER which soda tasted “better” OR which soda tasted “worse.” This particular participant was asked to indicate which soda tasted “better [worse].”

Figure 3. Stimulus from the form response/better condition in Study 4.

I think soda A tasted
better than soda B .

Participants then saw a completed form in which the participant had filled in (depending on condition) either that soda A tasted better than soda B, or that soda B tasted worse than soda A (see Figure 3). The design was therefore a 2 (response: free vs. form) * 2 (comparison: better vs. worse) factorial. If the effect of comparison framing in the previous studies resulted from inferences about the communicator's beliefs, making the choice of frame uninformative in the form conditions should attenuate or eliminate the effect.

After seeing the soda evaluations, participants were asked to rate how tasty they thought each soda was on 11-point scales (-5 = "Not at all tasty"; 0 = "Somewhat tasty"; 5 = "Extremely tasty"). They were then asked to recall which soda (A or B) the focus group participant had preferred.

We set a target sample size of 400 participants from Amazon Mechanical Turk and obtained complete responses from 405. We retained responses from the 338 participants who correctly recalled that soda A was preferred to soda B. We separately subjected participants' ratings of soda A and soda B to 2 (response format: free vs. form) * 2 (comparison: better vs. worse) between-subjects ANOVAs. For both sodas A and B, there was a main effect of comparison frame; soda A $F(1, 333) = 12.96, p < .001$; soda B $F(1, 333) = 83.17, p < .001$. There was no significant effect of response format; soda A $F(1, 333) = 3.21, p = .08$; soda B $F(1, 333) = .41, p = .52$. Nor, crucially, was there a significant frame * response format interaction; soda A $F(1, 333) = .70, p = .40$; soda B $F(1, 333) = .35, p = .55$. As in previous studies, the effect of comparison frame was stronger for soda B than soda A, $F(1, 335) = 46.28, p < .001$.² Means and standard deviations by condition are shown in Table 2.

² The test of this interaction was not significant (although it was directionally consistent) when including participants who failed the attention check.

Table 2. Means and standard deviations of soda quality ratings by condition in Study 4.

	Free Response		Form Response	
	“Worse” frame	“Better” frame	“Worse” frame	“Better” frame
	<i>M</i> (SD)	<i>M</i> (SD)	<i>M</i> (SD)	<i>M</i> (SD)
Soda A	1.26 (1.86)	1.76 (1.49)	1.43 (1.90)	2.24 (1.46)
Soda B	-2.37 (1.61)	-0.52 (1.49)	-2.37 (2.37)	-0.76 (1.45)

The results of Study 4 are incompatible with an inference-based alternative explanation of comparison framing effects. Even when it was clear that the communicator did not choose the comparison language, and thus there was no reason to believe that the choice of frame might reveal communicator beliefs, quality inferences were affected by framing. In fact, in the form conditions (where the choice of frame should be uninformative), the effects of comparison frame were sizeable, $d_{soda\ A} = .49$, $d_{soda\ B} = .85$.

We replicated these results in Study S5, which is reported in detail in the Supplemental Material. The design of this study was very similar to the current study except that comparison frame was manipulated within-participants: Participants saw both “better” and “worse” comparison frames (in counterbalanced order), either as free responses or pre-printed forms (between-participants, as in the current study). The “communicator inferences” account predicts that ratings of each soda should be more affected by framing when the comparative is a free response rather than a pre-printed form. However, the effect of framing was equally strong in either case (as shown by non-significant response format * comparison frame interactions), Soda A, $F(1, 171) = .81$, $p = .37$; Soda B, $F(1, 171) = .20$, $p = .66$.

It appears that explicit inferences about communicator beliefs cannot explain the effects of comparison frames on judgment and choice. It might be, however, that people implicitly make

inferences that they are not aware of or are unable to override. We return to this possibility and the idea of communicator intentions more generally in the General Discussion. We next turn to the implications of comparison framing for choice.

Study 5

In Studies 5 and 6, we further investigate the consequences of comparison framing for choice. We found some evidence in Studies 1a and 1b for such a consequence—when presented with a hypothetical choice between the inferior option and another randomly-chosen option, participants were more likely to prefer the randomly-chosen option under the “worse” comparison framing. This was the case even when the randomly-chosen option could logically be inferred to always be worse than the inferior option (e.g., in Study 1b, option B was in the top 10 but the randomly-chosen option was not). Showing consequences for choice is important for two reasons: First, changes in judgments that do not result in downstream behavioral changes may be inconsequential. Second, choices are not susceptible to the “scale confusion” alternative explanation described earlier (see Moore, 2007).

In the current study, participants were asked to imagine that they had the opportunity to try different candies and were shown a list of ten different fruit candy flavors. These were said to have been rated by a group of others from best to worst. Participants were then asked to imagine that a friend had provided some information about two of the candies: “Strawberry” (ranked #2) and “Orange” (ranked #3). Depending on condition, the friend either described Orange as significantly worse than Strawberry, or Strawberry as significantly better than Orange.

Participants were then asked to (hypothetically) choose 17 candies total from the full set of ten, and to indicate how many of each flavor they would choose (we chose 17 to prevent participants from choosing equal numbers of each candy; choices were forced to total to 17). We

were interested primarily in how many Strawberry and Orange candies people chose in each condition.

We requested 400 participants from Amazon Mechanical Turk and obtained complete responses from all 400. We retained for analysis the 341 participants (170 male, 169 female, 1 other/non-binary) who correctly recalled that their friend had preferred Strawberry to Orange. Of these, 174 had been randomly assigned to the “worse” condition and 167 to the “better” condition.

The number of candies chosen is a count, and values were not normally distributed (skewed and heavy-tailed; both $ps < .001$ by Stata’s “sktest” command). We therefore report results from negative binomial regressions (using Stata’s “nbreg”), but results are equivalent using t -tests. Participants chose the inferior candy (Orange) significantly more under the “better” comparison frame; $M_{better} = 1.66$, $SD_{better} = 1.68$, $M_{worse} = 1.07$, $SD_{worse} = 1.25$, $b_{frame} = .44$ [95% CI: .21, .68], $Z = 3.68$, $p < .001$. There was no significant effect of comparison frame on how many times the superior (Strawberry) candy was chosen, $M_{better} = 4.50$, $SD_{better} = 3.30$, $M_{worse} = 4.22$, $SD_{worse} = 3.40$, $b_{frame} = .06$ [95% CI: -.09, .22], $Z = .80$, $p = .42$. We tested whether this difference was itself significant in a mixed model, using the “menbreg” command in Stata. Candy (Orange vs. Strawberry) varied within participants and condition (better vs. worse comparison frame) varied between them. This model showed a significant interaction between candy and condition, $\chi^2(1) = 7.28$, $p = .007$, indicating that the effect of comparison framing was reliably stronger for the inferior candy.³

Discussion

³ This interaction was not significant in a repeated-measures ANOVA (i.e., assuming normality despite non-normal data), $p = .45$.

The results of the current study closely parallel the results when participants are asked for ratings rather than choices: there is an effect of comparison framing, but it is significantly stronger for the item described as inferior. In Study 6 we extend the effect of comparison framing to actual (i.e., incentivized) choices.

Study 6

In Study 6, we tested the effect of comparison framing on incentivized choices. We asked participants to make a choice between two movie rental vouchers. They could choose either a voucher for a film either from the iTunes “Top Rentals” section (described as containing “recently-released films that tend to be broadly popular” and cost up to 5.99 to rent) or from the iTunes the “99-cent Rentals” section (described as containing “less recent, more obscure, and less popular films”). All participants read that there were two films from the “Top Rentals” section: Film A and Film B. They further read either that according to the iTunes customer ratings, people thought that “Film A was better than Film B” or that people thought “Film B was worse than Film A.” All participants then read that films C-H were from the 99-cent rentals section; these films were not described further. Participants were then asked to choose whether they would like a voucher for a free rental of Film B or for a randomly-chosen film from C-H (the 99-cent rentals). They were told at the beginning of the study that 50 participants would be contacted and would actually receive what they had chosen.⁴

The sample size and analysis plan for this study were preregistered at <http://aspredicted.org/blind.php?x=6hk4iz>. All analytic choices in this study can be considered confirmatory unless otherwise specified. Following our preregistration, we requested 1,000 participants from Amazon Mechanical Turk and obtained complete responses from 1,006. We

⁴ iTunes does not actually allow one to gift rentals (only purchases), so all 50 randomly selected participants actually received a 5.99 bonus payment.

retained the responses of the 844 who correctly recalled that Film A was rated more positively than Film B. As predicted, participants were more likely to choose Film B over the 99-cent alternative under the “better” frame (320/437; 73%) than under the “worse” frame (257/426; 60%), $\chi^2(1) = 16.20, p < .001$. (The overall preference for Film B over the alternative is to be expected given the stated higher price of the “Top Rentals” films.)

Discussion

The results of Study 6 confirm the effect of comparison framing even for incentivized choices. By examining choices, rather than scale ratings, Studies 5 and 6 also speak against an important alternative interpretation of the previous studies; confusion about the response scale (Moore, 2007). Besides Study S1 in the appendix, the effects on choice found in Studies 5 and 6 further rule out this “scale confusion” alternative.

General Discussion

Across seven studies employing a variety of targets, comparison frames reliably influenced both judgment and choice. When comparisons were framed as “worse” (e.g., “B was worse than A”), both options, but particularly the inferior option, were evaluated more negatively than when comparisons were framed as “better” (e.g., “A was better than B”). According to the comparison framing account, these effects occur because of the different psychological effects of marked vs. unmarked comparisons. The marked comparison (“worse”) both evokes the negative end of the quality scale and commits the unfavorably-compared item to that end of the scale. Thus, this comparison frame leads to more negative evaluations of both items, but particularly the unfavorably-compared item (unless, as in Study 3, both items have already been committed to the negative end of the scale). The marked comparison frame also has specific effects on other items from the same set—i.e., those at the bottom of the quality range are seen as worse (Study

2). Finally, the effects of comparison frames go beyond evaluations to choices: Under the “worse” frame, the unfavorably-compared item was chosen less vs. alternative options (Studies 5 and 6).

Relation to previous research

We will next discuss how our studies relate to previous research and existing theories of framing.

Comparative judgment. There is a substantial body of research specifically examining comparative judgments. The most prominent model of comparative judgment is Giladi and Klar’s “LOGE” model (Giladi & Klar, 2002; Klar & Giladi, 1997, 1999). According to the LOGE model, when people make comparative judgments (e.g., “How does this student compare to others in his class?”) they often over-weight their evaluations of the focal target and underweight their evaluations of the reference group. So, for example, in a class of excellent students, all might be rated as above-average relative to the class; conversely, in a class of under-performers, all might be rated as below-average (again, relative to the class). This might seem similar to the comparison framing effects we document here, but there are substantial differences. The LOGE model specifically concerns the contamination of relative judgments by characteristics of the target, whereas comparison framing entails changes in *target* evaluations resulting from differences in how relative evaluations (i.e., comparisons) are expressed. Additionally, some of the empirical support for the LOGE model may actually result from participants being confused about what question they are being asked (i.e., they actually think they are being asked for a judgment of the target, rather than a relative judgment; Moore, 2007). Studies 5 and 6 speak against a “scale confusion” account of comparison framing. The relationship between the LOGE model and our findings is therefore more a general similarity in

approach (e.g., normative judgment errors related to comparative judgments) rather than competing explanations of the same empirical facts.

Attribute framing. The current findings are most closely related to previous research on framing—specifically, attribute framing (Levin & Gaeth, 1988; Levin, Schneider, & Gaeth, 1998). Like attribute frames, comparison frames are different but logically equivalent descriptions of the same facts that lead to large differences in evaluations. However, there are also important distinctions. Most obviously, a comparison is not an “attribute” (e.g., fat content) in the way that the term has been used in the literature. As well, attribute framing is understood as affecting a single target item, whereas the effect of comparison frames extends not only to the items being compared but also to other members of the set from which they were drawn. Most importantly, the most widely-accepted account proposed to explain the effects of attribute frames cannot explain all the current results. According to this account, attribute framing effects result from automatic valenced associations evoked by the frames. A negative frame (“75% fat”) evokes other negative information in memory, whereas a positive frame (“25% lean”) evokes other positive information. This information then biases evaluations of the target of judgment (Levin, Schneider, and Gaeth, 1998). For the current results, this account would hold that the “better” comparison frame activates positive associations, whereas the “worse” frame evokes negative associations; these would then carry over to the items being compared. The valenced associations account is thus compatible with some of the effects of comparison frames. However, it does not predict an asymmetry such that comparison frames have stronger effects for the inferior option, or the conditions under which this asymmetry should not be observed (that is, when both items are known to be of low quality). Nor does it predict effects of framing for other

items in the set at the bottom of the quality scale. Thus, the valenced associations account cannot explain the entirety of the results we observed.

A recent account of attribute framing proposes that participants focus on the presented frame and neglect the complementary frame, leading to downstream changes in judgment (Gamliel & Kreiner, 2020; Kreiner & Gamliel, 2018). Although there are important differences between the attentional and valenced associations account, for our purposes they are irrelevant. Like the valenced associations account, an attentional account would not predict asymmetric effects of comparison frames, nor would it predict effects for other low-quality items in the set.

Leakage of information. A very different interpretation of attribute framing effects is that they may result from the normatively-appropriate use of information “leaked” by the communicator’s choice of frame (Leong, McKenzie, Sher, & Müller-Trede, 2017; McKenzie & Nelson, 2003; Sher & McKenzie, 2006, 2008). Broadly, this account posits that people go beyond what is stated explicitly to infer information from a communicator’s utterances, and that this can be rational. Under this account, then, framing effects can result from listener’s normatively-appropriate use of information inherent in the choice of frame. This account is quite specific about how the choice of frame conveys information. Specifically, the hypothesis is that speakers tend to choose the frame “that has increased relative to the reference point.” For example, speakers would be more likely to describe ground beef as “25% fat” (vs. “75% lean”) if 25% was an unusually high fat percentage (McKenzie & Nelson, 2003). As we described in the Introduction, one could modify this hypothesized mechanism to apply to both items being compared (e.g., “both the items being compared are below average” or “both items being compared are above average”). In this case, the prediction would be that comparison frames

should affect evaluations of the superior and inferior items equally (i.e., it would not predict asymmetric inferences). Neither would it predict effects for any other items in the choice set.

Inferences about communicator beliefs. The leakage of information account has sometimes been construed much more broadly to hold that people go beyond the literal content of the information given to make inferences based on their background knowledge of how speakers typically use language. This might even occur automatically, such that even normatively uninformative changes in language (as in our Study 3) might give rise to inferences about speaker intent that would change judgments (for an argument along these lines, see Sher & McKenzie, 2006). In this broader sense the leakage of information account is compatible with our results. However, the claim that people go beyond the literal content of an utterance to infer speaker meaning is not unique to the leakage information account, nor is does it make specific predictions about what inferences will be made.

It is obviously true that listeners often interpret speakers non-literally. For example, most people automatically interpret “Can you turn off the light?” as a request, not as a question about their abilities. This observation was first formalized by Grice (1975), who coined the term “conversational implicature” to describe the (inferred) meaning of an utterance, as distinct from its literal content. However, this general observation does not specify *how* people go beyond the literal content to infer meaning (Grice went on to describe “conversational maxims” of how listeners might interpret utterances non-literally). In the case of comparison frames, the observation that laypeople make inferences from the language used in comparisons does not make testable predictions without the additional specification of *what* meaning listeners will extract from different comparison frames, and why. In other words, positing that people automatically extract information implicit in a speaker’s utterances does not explain what

information might be encoded by speakers, what the encoding rules are, what listeners will extract, and how this will affect inferences. The markedness account supplies these details. Likewise, the leakage of information account posits a specific rule by which listeners extract information from a speaker's choice of frame—i.e., that listeners assume speakers choose a frame that exceeds a reference point (for example, describing meat as “25% fat” rather than “75% lean” if most meat is less fatty). The two accounts, therefore, describe very different ways in which listeners extract information from speaker utterances—even though both accounts posit that listeners go beyond the logical content of the utterance in doing so.

Generality of comparison framing effects

We found effects of comparison frames for a wide range of targets (including consumer goods, personnel, food and drink, and films), across different response scales (Likert ratings, percentiles, and choices), and for consequential (i.e., incentivized) choice. Although robust across targets and scales, the effect of comparison frames depends on the distinction between marked and unmarked adjectives in English, and we would not expect languages without this distinction to produce comparison framing effects. However, most languages do seem to mark “worse” in the same way English does—Cruse (1992) found that “better” and “worse” are (un)marked as in English, French, Turkish, and Chinese—the only exception in the languages he examined was Macedonian.

Additionally, the results of Study 1a and Study 2 show that the psychological process underlying inferences from marked comparatives doesn't require the specific antonyms “better” and “worse.” In Studies 1a and 2 (and their replications), the comparisons were “significantly more qualified” and “significantly less qualified.” These comparatives evidently encode scale positions (and, in the case of “less qualified,” commit items to the negative end of the scale) just

as “better” and “worse” do. Future research should examine the psychology of these comparatives in more detail. For example, is “less qualified” implicitly translated to “(more) unqualified,” or does it directly evoke the negative end of the qualification scale? Does the same psychology underlie the processing of other similarly phrased comparisons (e.g., is “significantly less intelligent” equivalent to “unintelligent”)? It may in fact be that many or most negative evaluative comparisons have psychological effects similar to “worse” in that they entail scale position and asymmetric commitment.

Conclusion

We observe robust and general effects of logically-equivalent comparison frames on people’s judgments and choices. These effects are asymmetric, affecting inferior items more strongly than superior ones, and extend to other items beyond those being directly compared. It is common that different verbal descriptions of the same problem lead to different judgments and choices (e.g., Fox & Rottenstreich, 2003; Shafir, 1993; Tversky & Kahneman, 1981). The current results demonstrate that understanding how language is used can be essential in understanding how this occurs.

References

- Chase, W. G., & Clark, H. H. (1972). Mental operations in the comparison of sentences and pictures. In *Cognition in Learning and Memory* (pp. vii, 263–vii, 263). John Wiley & Sons.
- Clark, H. H. (1969a). Linguistic processes in deductive reasoning. *Psychological Review*, 76, 387-404.
- Clark, H. H. (1969b). Influence of language on solving three-term series problems. *Journal of Experimental Psychology*, 82, 205–215.
- Clark, H. H., & Chase, W. G. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology*, 3, 472–517.
- Cruse, D. A. (1976). Three classes of antonym in English. *Lingua*, 38, 281–292.
- Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, 14, 195–200.
- Gamliel, E., & Kreiner, H. (2020). Applying fuzzy-trace theory to attribute-framing bias: Gist and verbatim representations of quantitative information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46, 497–506.
- Giladi, E. E., & Klar, Y. (2002). When standards are wide of the mark: Nonselective superiority and inferiority biases in comparative judgments of objects and concepts. *Journal of Experimental Psychology: General*, 131, 538-551.
- Gleitman, L. R., Gleitman, H., Miller, C., & Ostrin, R. (1996). Similar, and similar concepts. *Cognition*, 58, 321–376.
- Greenberg, J. H. (1966). Language universals. *The Hague: Mouton*.

- Harris, R. J. (1973). Answering questions containing marked and unmarked adjectives and adverbs. *Journal of Experimental Psychology*, 97, 399–401.
- Klar, Y., & Giladi, E. E. (1997). No one in my group can be below the group's average: A robust positivity bias in favor of anonymous peers. *Journal of Personality and Social Psychology*, 73, 885-901.
- Klar, Y., & Giladi, E. E. (1999). Are most people happier than their peers, or are they just happy?. *Personality and Social Psychology Bulletin*, 25, 586-595.
- Kreiner, H., & Gamliel, E. (2018). The role of attention in attribute framing. *Journal of Behavioral Decision Making*, 31, 392–401.
- Lehrer, A. (1985). Markedness and Antonymy. *Journal of Linguistics*, 21(2), 397–429.
- Leong, L. M., McKenzie, C. R., Sher, S., & Müller-Trede, J. (2017). The role of inference in attribute framing effects. *Journal of Behavioral Decision Making*, 30, 1147-1156.
- Levin, I. P. (1987). Associative effects of information framing. *Bulletin of the Psychonomic Society*, 25, 85–86.
- 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, 374-378.
- 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, 149-188.
- McKenzie, C. R. M., & Nelson, J. D. (2003). What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin and Review*, 10, 596–602.

- Moore, D. A. (2007). When good = better than average. *Judgment and Decision Making*, 2, 277-291.
- Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory & Cognition*, 21, 546–556.
- Sher, S., & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition*, 101, 467–494.
- Sher, S., & McKenzie, C. R. M. (2008). Framing effects and rationality. In M. Oaksford, & N. Chater (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 79–96). Oxford: Oxford University Press.
- Talmy, L. (1978). Figure and ground in complex sentences. In J. Greenburg, C. Ferguson, & M. Moravcsik (Eds.), *Universals of human language* (IV). Stanford: Stanford University Press.
- Trabasso, T., Rollins, H., & Shaughnessy, E. (1971). Storage and verification stages in processing concepts. *Cognitive Psychology*, 2, 239–289.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453-458.