Baby, I Can See Your Halo[1]: How the Halo Effect, Loss Aversion, and Strategic Attribute Disclosure Neglect Can Inform Negative Advertising

Madison Sultan

Advised by David Owens

May 2019

Abstract

This paper analyzes how the halo effect, loss aversion, and neglect of the strategic nature of information revelation are present in consumer's sequential learning and belief updating about the quality of products. The implications of these biases are analyzed to determine the circumstances under which producers will optimally engage in negative advertising. A theoretical model is developed and tested using an experiment conducted on Prolific Academic using data from Consumer Reports.

---

[1] Lyrics from Beyoncé's "Halo"

1. Introduction

The practice of negative political campaigning, or emphasizing the negative traits of competitors versus positive traits about oneself, has been a popular political campaign tactic for decades (Fridkin & Kenney, 2004) and the popularity has increased recently. A study done at Wesleyan University found that 48% of all campaign advertisements for the 2018 election cycle were negative with a 61% increase in the number of negative advertisements compared to the midterm elections in 2014 (Rubenstein, 2018). Numerous studies have examined this trend and how it affects various political outcomes such as voter turnout and the probability of being elected (Barton, Castillo, & Petrie, 2016; Fridkin & Kenney, 2004; Landi, 2004; Lau, Sigelman, & Rovner, 2015). From this research, the effects of negative campaigning are well documented.

In comparison, the use of negative advertising, the equivalent of negative campaigning in consumer goods, has been studied much less. However, this technique is still used in marketing. For instance, during the 2019 Superbowl, Bud Light gained publicity for a series of television advertisements and accompanying billboards which simply claimed "Coors Light uses corn syrup" ("MillerCoors sues Anheuser-Busch over corn syrup ads," 2019). Given the controversy this sparked, an academic study which isolates the potentially relevant variables can help to better understand why a brand this prominent might have the incentive to use this advertising strategy.

This paper will begin to explore this question by experimentally comparing the effectiveness of revealing information about the quality of one's own product traits versus information about the traits of a competitor's product. Additionally, it incorporates a theoretical model and experiment which test how consumers re-evaluate their judgements about products as they learn new information about how its various traits have been rated. The evidence found supports hypotheses that consumers tend to overestimate the correlation between a products' traits and that they fail to understand the strategy behind a firm's use of negative advertising.

2. Literature Review

a. Halo Effect

The halo effect is a phenomenon where people tend to think that the separate traits that characterize a person or an object are more related than they are in actuality. First coined by Edward Thorndike, it has traditionally been studied in the field of psychology when raters make subjective judgements about the personality, attributes, and general merits of other people. Thorndike found that they tended to "think of the person in general as rather good or rather inferior and to color the judgements of the qualities by this general feeling" (Thorndike, 1920, p. 25). This theory involves the overall quality or some other particularly prominent attribute of a subject influencing the processing of information about, and subsequent judgments of, their other separate traits. For instance, a number of papers have found that when people have been pre-tested to be attractive, they are rated as more intelligent when they produce identical work to those who are unattractive (Landy & Sigall, 1974), they are perceived to be better political candidates (Verhulst, Lodge, & Lavine, 2010), and they are expected to have more social personalities and live better lives (Dion, Berscheid, & Walster, 1972).

However, there is a clear potential issue to these measures of the halo effect—what if these traits are in fact highly related? When the method of analysis is to simply measure the correlation amongst ratings, it is unclear if there is truly an error being made or if this correlation in ratings is an accurate reflection of the correlation amongst the traits. Later literature recognized this problem and began to differentiate between the true relationship amongst traits, known as the valid halo, and the extent to which the ratings of the traits are correlated beyond the true relationship, or the invalid halo (Cooper, 1981).

To differentiate between these, researchers often will experimentally manipulate traits to be independent of one another. For instance, in one study, male respondents were asked to rate the quality of an essay. Before reading it, they were given a fabricated biography of the author which included a photo. Some respondents received a photo of a women who was pre-rated to be attractive, some of a woman who was pre-rated as unattractive, and some received no photo at all (but still received the biography). Despite reading *identical* essays, those who were shown an attractive author rated the essay as significantly better than those who received no photo, and those shown an unattractive author rated the author as significantly worse. In this case, despite the author's attractiveness being completely independent of the quality of the essay, it still had a significant effect on this rating (Landy & Sigall, 1974).

However, just as was true in the earlier examples, this experiment also was examining how one subjective trait judgement (attractiveness of the author), influenced another subjective trait judgement (quality of the essay). This paper will strengthen the literature on the halo effect by examining how the knowledge of the quality of a trait that has been objectively rated influences expectations about the quality of another objectively rated trait. Additionally, most of the literature has focused on how the perception of one trait influences the perception of another, and this paper will examine how the knowledge of one trait influences *expectations* of an unknown trait. More specifically, it predicts that consumers will systematically overestimate the value of unknown traits for products that perform well in known traits and underestimate them for those that perform poorly.

b. Reference-Dependence and Loss Aversion

If the halo effect occurs, then following the revelation of the quality of one of a product's traits, consumers may form biased expectations about the quality of its other traits. If the true values of these traits are later revealed, the relationship between these expectations and the revealed value could influence perceptions of the product overall as theorized by Daniel Kahneman and Amos Tversky's Prospect Theory.

This model proposes that when consumers evaluate a choice, they may not only care about the final information they have about this option, but also about the expectations they had throughout the process of learning about it. For instance, consider someone choosing which college to attend. Based on their knowledge of a school's extremely low acceptance rate, suppose that they expected that college would have the best Economics department in the country. If they later discovered that it actually had the second-best department, they might be disappointed in this information and might be less likely to attend the school, even if they would have been more than happy to attend the second-best department if this expectation had not been set. In short, options are often considered subject to a reference point (Kahneman & Tversky, 1979) which can be determined by many different things such as the endowment of peers or a personal former endowment, but this paper involves reference points formed based on expectations of the value. Existing literature has shown that expectations do serve as reference points when considering prices (Kalwani, Yim, Rinne, & Sugita, 1990) and the quality of an item (Hardie, Johnson, Peter, & Fader, 2007) as proposed in the seminal paper "A Model of Reference-Dependent Preferences" (Kőszegi & Rabin, 2006).

Prospect theory also proposes that when someone has more of an asset than their reference point, they are in the "domain of gains" and when they have less than this value, they are in the "domain of losses. Movement into the domain of losses decreases expected utility and movement into the domain of gains decreases expected utility. Additionally, the impact is greater in absolute value for movement into the domain of losses than for movement into the domain of gains (Kahneman & Tversky, 1979). This paper uses reference-dependence and loss aversion and theorizes that when information is revealed that does not meet expectations, the negative impact is greater in magnitude than the positive impact of information that exceeds expectations by an equivalent amount. This provides some motivation for why a firm may want to engage in negative advertising to disparage their competitor as opposed to positive advertising about their own product.

a. Strategic Attribute Disclosure Neglect

Finally, when firms decide what information to highlight, they are strategically disclosing some details while strategically withholding others. Therefore, when presented with information, consumers can make inferences about the unknown information by examining what the firm's objective is for this disclosure. However, numerous experiments have found that consumers ignore the selection bias present in strategically disclosed information (Brown & Fragiadakis, 2018; Deversi, Ispano, & Schwardmann, 2018; Koehler & Mercer, 2009).

This paper adds to the literature by examining if this bias occurs when consumers are aware that the producer could have revealed information about themselves or their competitor. It theorizes that similarly to the above scenarios, consumers will not properly factor in the method by which they learned information and will therefore neglect to learn from the information that was strategically withheld from them. This literature along with that on the halo effect, reference-dependent preferences, and loss aversion develop the following model.

3. Theoretical Framework

a. Overview

This game examines a situation where there are two products, X and Y, that are characterized by two separable traits, A and B, and each of the products' traits have been randomly assigned a true quality rating by nature. The firms that produce X and Y know the true value of A and B for both products. Consumers have prior beliefs about the distribution of these ratings for the entire category of products that X and Y are a part of, but they do not initially know any differentiating information about the ratings of X and Y specifically. In stage 1, the true value of A for both X and Y is then revealed to the consumer, and they update their beliefs about the rating of B based on their impression of the correlation between the quality ratings of A and B. They also form opinions about the overall quality of the products. In stage 2, suppose the producer of product X knows the true rating of B for both X and Y and is then given the opportunity to reveal one of these values to the consumer. Since the overall evaluations of the products depend on the consumer's knowledge and expectations of the products' traits, the producer will reveal whichever trait value maximizes the probability that the overall evaluation of X will be greater than the overall evaluation of Y. Following this revelation, the consumer then re-evaluates their beliefs about the value of b for the unrevealed product and then produces an overall evaluation of both products.

b. Halo Effect

Let the consumer's prior belief about the distribution for trait A for the entire category of products be $f_A(a)$ and their prior belief about the distribution for trait B be $f_B(b)$. Therefore, their prior expectation is that any given product in that category has a true value of A equal to $E(f_A(a))$ and a true value of B equal to $E(f_B(b))$. Then, following the revelation that for product

X, the true value of A is a_X, the consumers will update their beliefs about the true quality of B following Bayes' posterior probability model (Larsen & Marx, 2017):

$$f_{B|A=a_X}(b) = \frac{f_B(b_X) * f_{A|B=b_X}(a_X)}{\int_{-\infty}^{\infty} f_A(a_X)} \quad (1)$$

However, if consumers are biased by the halo effect, then they will believe that the traits are more closely related than is true in actuality. Therefore, in their probability estimates, they are overweighting the connection between traits, $f_{A|B=b_X}(a_X)$ and underweighting their prior beliefs about the distribution of B alone, $f_B(b_X)$. As such, their biased beliefs, $f_{B|\widehat{A=a_X}}(b)$, about the quality of B following the revelation of a_X is:

$$f_{B|\widehat{A=a_X}}(b) = \frac{f_B(b_X) * g(f_{A|B=b_X}(a_X))}{\int_{-\infty}^{\infty} f_A(a_X)} \quad (2)$$

where g(·) is some function that inflates $f_{A|B=b_X}$ when a_X is relatively high and deflates $f_{A|B=b_X}$ when a_X is relatively low. Therefore, their estimations about the quality of B will be biased towards the information they learned about the quality of A.

c. Reference-Dependence and Loss Aversion

If consumer's exhibit reference-dependent preferences, then their overall evaluations of the products will not only be affected by their knowledge about the traits that characterize it, but also by how this information compares to their expectations of it. Additionally, loss aversion predicts that the disappointment of information falling below expectations will have a greater effect on the overall rating in absolute value than when information is pleasantly above expectations. Therefore, their overall rating for product X in stage 1 can be modeled as:

$$OverallRating_1^X = a_X + E(b_X)_1 + v(a_X|E(a_X)_0) + v(E(b_X)_1|E(b_X)_0) \quad (3)$$

where $v(i|j)$ is the generic function defined as

$$v(i|j) = \begin{cases} \ln((i-j)+1)\ for\ i \geq j \\ -\eta(\ln(-(i-j)+1))\ for\ i \leq j \end{cases} \quad (4)$$

Here, the subscript 0 denotes the expectations of that trait in the prior stage. Since at that point, the consumer had no differentiating information about X or Y, these values are theoretically equal to the mean of their beliefs about the rating of that trait for the entire category of products, or E($f_A(a)$) and E($f_B(b)$). Additionally, $v(i|j)$ denotes the reference-dependence value of how a change in expectations from j to i may affect overall evaluations. The term η represents how surprisingly bad information compares to surprisingly good information. If a consumer is loss averse, they will weight decreases in expectations more strongly than increases, and thus, this term will be greater than 1. An identical model can also be used to examine their overall ratings for product Y.

If consumers exhibit the halo effect and also have this reference-dependent term in their preferences, then the revelation of A has an even greater impact on evaluations of the overall quality. Without the halo effect, $E(b_X)_1$ will be updated solely by the true amount that A and B are related, and the difference between $E(b_X)_1$ and $E(b_X)_0$ will reflect this more modest change. However, with the halo effect, consumers more drastically change their expectations of B following the revelation of information about A which broadens the difference between $E(b_X)_1$ and $E(b_X)_0$.

In the second stage, consumers are then learning new information about the true quality of B for either X or for Y. Therefore, the model for the consumer's overall rating of product X in stage 2 is the following:

$$OverallRating_2^X = a_X + E(b_X)_2 + v(a_X|E(a_X)_0) + v(E(b_X)_1|E(b_X)_0) + v(E(b_X)_2|E(b_X)_1) \quad (5)$$

where $v(i|j)$ is defined as described in (4)

Again, the same model can also be used to describe the overall ratings of Y. This model reflects their new expectations about the quality of B, $E(b_X)_2$. This term is an expectation if the producer decided to reveal the true quality of B for the other product, b_Y, or it is the actual value b_X if this value was revealed. Additionally, compared to the overall ratings in stage 1, there is a new term $v(E(b_X)_2|E(b_X)_1)$ which reflects how this expectation/knowledge of the true value of B compares to the expectations of it in the previous stage.

d. Self versus Other Reveal

$v(E(b_X)_2|E(b_X)_1)$ is of particular interest to the producer in their decision about whether to reveal b_X or b_Y. If consumers had not exhibited the halo effect, then this term should be 0 in expectation for both products since their stage 1 estimations would be accurate on average. However, the halo effect means that products that were of a lower quality than was expected on trait A have falsely low expectations for trait B in stage 1 ($E(b_X)_1 < E(b_X)_2$ in expectation) and products that did better than expectation for trait A have inflated expectations for B ($E(b_X)_1 > E(b_X)_2$ in expectation). Additionally, with the halo effect, $v(E(b_X)_2|E(b_X)_1)$ will be negative for products that did well in trait A and positive for products that did poorly. First, consider the situation in which either b_X or b_Y is randomly revealed, and as such, the revelation of one should not result in a change in expectation for the other.

- If both a_X or a_Y are equal and are less than $\mathrm{E}(f_A(a))$, or both did equally poorly compared to expectations, then the revelation of B for either product will result in this term being positive. It would therefore increase $OverallRating_2^X - OverallRating_2^Y$ most if b_X is revealed since it increases $OverallRating_2^X$ without affecting $OverallRating_2^Y$.
- If $\mathrm{E}\big(f_A(a)\big) - a_X = a_Y - \mathrm{E}\big(f_A(a)\big)$, or if X did as poorly compared to expectations as Y did well, then $v(E(b_X)_2|E(b_X)_1)$ would be positive and $v(E(b_Y)_2|E(b_Y)_1)$ would be negative. Therefore, since $v(E(b_Y)_2|E(b_Y)_1)$ would be more negative than

$v(E(b_X)_2|E(b_X)_1)$ would be positive, the overall rating differential is increased most if b_Y is revealed.

- If $a_X - \mathrm{E}\big(f_A(a)\big) = \mathrm{E}\big(f_A(a)\big) - a_Y$, or if X did as well compared to expectations as Y did poorly, then $v(E(b_X)_2|E(b_X)_1)$ would be negative and $v(E(b_Y)_2|E(b_Y)_1)$ would be positive. Since the information about b_X hurts expectations of X more than b_Y helps Y, the differential is diminished less when b_Y is revealed.
- If a_X or a_Y are equal and are greater than $\mathrm{E}(f_A(a))$, or if both did equally well, then both $v(E(b_X)_2|E(b_X)_1)$ and $v(E(b_Y)_2|E(b_Y)_1)$ would be negative in expectation, so revealing b_Y increases the differential while revealing b_X decreases it. Therefore, it is more beneficial to the relative overall rating of X if b_Y is revealed.

e. Strategic Attribute Disclosure Neglect

However, these predictions change if consumers are aware of the strategic nature of the information that has been revealed to them and they properly respond. Suppose that the consumer is aware that the producer of X had the opportunity to disclose either b_X or b_Y with the intention of increasing the probability that $OverallRating_2^X > OverallRating_2^Y$ so that the consumer would choose product X over product Y. If they are strategically shown b_X, compared to someone who was randomly shown this information, their estimates of b_Y should be higher since the worse that b_Y is, the more likely a strategic producer of X would have chosen to reveal its value. Similarly, if the consumer is strategically shown the value of b_Y, they should have lower estimates of b_X compared to someone who received this information randomly because the better that b_X is, the more likely it would've been that the producer would have shown this value. Reexamining each of the above cases:

- In the case where both a_X or a_Y are equal and are less than $\mathrm{E}(f_A(a))$, b_X was revealed to increase $OverallRating_2^X$. However, if the consumers see this and adjust up their estimates of b_Y, this would increase $OverallRating_2^Y$, potentially counteracting the positive affects of the revelation of b_X.
- When $\mathrm{E}\big(f_A(a)\big) - a_X = a_Y - \mathrm{E}\big(f_A(a)\big)$, b_Y was revealed since the decrease to $OverallRating_2^Y$ outweighed the potential increase to $OverallRating_2^X$. Yet if this reveal is correctly responded to, the consumer would lower their ratings of b_X and as such would lower $OverallRating_2^X$, again potentially counteracting the benefits of revelation.
- When $a_X - \mathrm{E}\big(f_A(a)\big) = \mathrm{E}\big(f_A(a)\big) - a_Y$, b_Y was revealed since this would raise $OverallRating_2^Y$ less than revealing b_X would decrease $OverallRating_2^X$. Yet if the consumer recognizes the strategy, they will decrease their expectations about b_X, decreasing $OverallRating_2^X$, and potentially decreasing the differential even more.
- When both a_X or a_Y are equal and are greater than $\mathrm{E}(f_A(a))$, b_Y was revealed as it decreased $OverallRating_2^Y$. Again though, if the consumer lowers their expectations for b_X, this will lower $OverallRating_2^X$ counteracting the positive effects of the reveal.

In each of these cases, if consumers are properly responding to the strategic nature of the information they receive, this will attenuate the producer of X's benefits of being able to choose what information to reveal. However, many consumers may neglect the nature in which information is revealed to them when they are incorporating new signals about the products' quality. Therefore, they will not properly adjust their estimates of b_Y when they are revealed b_X, and they will not properly adjust their estimates of b_X when they are revealed b_Y. As such, the first set of predictions about the effectiveness of revealing b_X or b_Y hold. To test these theoretical predictions, the following experiment was designed.

4. Experimental Design

a. Strategy

First, data needed to be found that included unbiased ratings of products. These products needed to have ratings for two traits that can be separately analyzed and in conjunction, almost entirely categorize the product. To test for the halo effect, the ratings for one of these traits needed to be revealed to consumers and then their expectations about the rating of the other trait could be collected. The relative relationship between the revealed trait and the expectations could be compared to the actual relationship between the true ratings. For loss aversion, consumers needed to form expectations about the quality of a trait and these expectations needed to be elicited. Then, the effect on overall ratings could be compared for consumers who learned information that exceeded their expectations versus fell below their expectations. To examine neglect of strategic disclosure, some consumers needed to be shown the second trait in an explicitly strategic way and some needed to be shown it in an explicitly random way. Then, expectations about the unrevealed trait value could be compared across the strategic and random disclosure groups.

b. Consumer Reports Data

Product data were collected from Consumer Reports because this agency is the best in industry for unbiased product reviews (Versability, 2018; "What We Do - Consumer Reports," n.d.). Additionally, the site has overall ratings of products as well as ratings for different traits that categorize the products. For this experiment, data on the ratings for 24 different camcorders were used. Camcorders were chosen as they have two distinct and important traits that categorize each item—their audio and video quality. The frequency of these ratings can be seen in the following figures.

Figures 1-3, Table 1: Camcorder Ratings

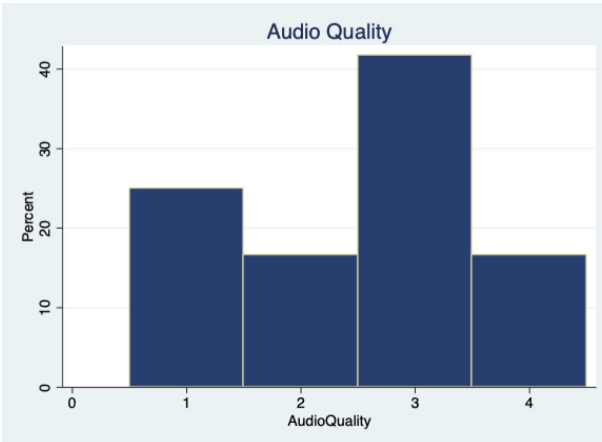

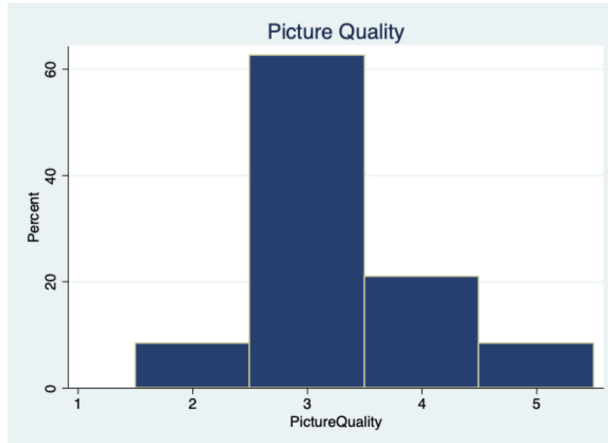

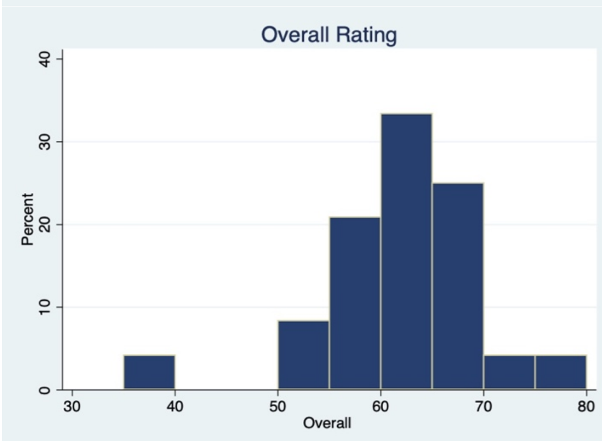

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Overall	24	61.67	8.014	36	77
PicQuality	24	3.292	0.751	2	5
AudQuality	24	2.500	1.063	1	4

Note that despite the fact that each trait was rated on a 1-5 scale, no products received a 5 for audio quality, and no products received a 1 for picture quality.

Table 2 examines the correlation between each of the ratings. The correlation between the trait ratings and the overall quality rating are each significant at the 1% level indicating that the individual traits are in fact highly related to the overall rating of the product. Additionally, the correlation between each of the traits and the overall rating is very similar indicating that each are equally important to the overall quality of the product. The correlation between the traits is insignificant indicating that knowledge of the rating of one trait does not provide any real information about the rating of the other trait for any given product.

Table 2: Rating Correlations

	Overall	Picture Quality	Audio Quality
Overall	1		
Picture Quality	0.5806*** (0.0029)	1	
Audio Quality	0.5664*** (0.0039)	0.1907 (0.3721)	1
n=24, * p < 0.1, ** p < 0.05, *** p < 0.01.			

c. Methodology

The experiment was conducted using a survey created on Qualtrics. The survey link was posted on the distribution platform, Prolific Academic, and 200 participants were recruited. When participants entered the survey, they were told that their task would include making estimations about how various camcorders had been rated on Consumer Reports. They were told that each camcorder's audio and picture quality is rated on a discrete scale of 1-5, and its overall quality is rated on a discrete scale of 1-100. The participants were also given a frequency distribution for how the camcorders had been rated for each of the two traits, and they were given a histogram of the overall quality ratings which can be seen in the above histograms. These distributions serve as the prior distributions from which participants could begin forming their estimations about how an individual product might be rated. They were also told that in addition to $1 in base pay, for each of the 7 estimations they would make, they could earn an additional 10 cents which would be rewarded based on the accuracy of their predictions using the incentive-compatible Quadratic Scoring Rule (Brier, 1950)[2].

On the next screen, two of the camcorders were chosen as their option set, one of which was labeled X and the other was Y. The participants were told the audio quality rating for both

[2] See appendix for a full description of payment.

products and then for each, they were asked to indicate the probability with which they thought that the picture quality had received each of the possible ratings, 1, 2, 3, 4, or 5. They could indicate anywhere from a 0 to 100% probability that it had been rated each individual rating, but they could not continue until the probability for the 5 ratings totaled 100[3]. They went through the same process to indicate their beliefs about the overall quality ratings, but instead they indicated the probability that the rating fell within a given 5 point range. Each of these ratings were stored as their stage 1 estimations[4].

Before stage 2, the participants were randomly assigned to either the random condition or the strategic condition. In the random condition, participants were shown, “An algorithm was given the ability to randomly reveal either the picture quality rating for X or for Y. It has randomly chosen to reveal that product X/Y received a __ for its picture quality rating”, and it showed the picture quality rating for either X or Y with equal probability. They then indicated their stage 2 beliefs about the picture quality rating for the product that had not been shown to them, as well as their new beliefs about the overall quality ratings for both items.

In the strategic condition, participants read, “An algorithm designed to increase your evaluation of X compared to your evaluation of Y was given the option to either reveal the picture quality rating for X or for Y. It has chosen the reveal that product X/Y received a __ for its picture quality rating.” The rating for X was revealed if and only if the true picture quality rating was above the expected value for this rating indicated in stage 1 by more than the true picture quality rating for Y was *below* expectations of it, or if:

[3] For example, they might indicate that it received a 1 with 5% probability, 2 with 15%, 3 with 50%, 4 with 20%, and a 5 with 10% probability.
[4] One thing to note with this methodology is that unfortunately, just as for the individual trait ratings, participants were making inferences about what the overall rating of the product might be as opposed to providing subjective impressions about which product they would prefer. This modification was necessary to abide by Prolific Academic’s terms of use. Issues with this are discussed further in the limitations section.

$$PicQualX - ExpPicQualX > ExpPicQualY - PicQualY \quad (6)$$

In other words, they were revealed the value of X only if this value was more surprisingly good than the value of Y was surprisingly bad. These participants also indicated their beliefs about the picture quality rating of the product they were not shown, as well as about the overall quality ratings for both products.

5. Data, Hypotheses, and Statistical Methodology

   a. Halo Effect

The halo effect is present in this experiment if participants use the information they have on the quality of one trait as too diagnostic of the quality of the other trait. In this case, this means that the relationship between the rating for a trait they know, audio quality, and their expectations about the rating for the trait they do not know the value of, picture quality, will be significantly stronger than the true relationship between these two traits.

To examine this, two regressions were run for both product X and product Y. In the first, the true audio quality was used as a predictor for the participant's expected value for the picture quality as indicated through their stage 1 probability estimate. In the other, the true audio quality was used as a predictor for the true picture quality. The following show the regressions run for product X, but identical regressions were run using the data for product Y.

$$XStg1ExpPic_i = \alpha_i + \beta_1 XAudRating_i \quad (7)$$

$$XPicRating_i = \alpha_i + \beta_2 XAudRating_i \quad (8)$$

Table 3: Halo Effect Variables

Variable	Meaning
$XStg1ExpPic_i$	The expected value for the picture quality rating of product X (although a similar variable exists for Y) as indicated by the participant in stage 1, following the reveal of the audio quality.
$XAudRating_i$	The true audio quality rating. This rating has just been revealed to the participant.
$XPicRating_i$	The true picture quality rating. This is currently unknown to the participant.

If the halo effect is present in the study, then a χ^2 test will show that the coefficient on $XAudRating_i$ is statistically significantly different in the two regressions.

b. Reference-Dependence

If participants exhibit reference-dependent preferences, then their overall ratings in each stage will be affected by both their current expectations/knowledge about the quality of each of the trait ratings but also by how these compare to their expectations about the quality in previous stages. Therefore, when the following regressions are done (again for both X and Y) in stages 1 and 2, the coefficients on the reference-dependent variables will be significant. Note that, in addition to doing the regression separately for products X and Y, it was also separated based on whether the participant had been revealed the picture quality rating for X or for Y. This is to allow for the fact that respondents might respond differently to a change relative to their expectations when the change is due to explicit knowledge (they know for certain that the picture quality is some value) versus just a change in expectations.

$$XStg1ExpOv_i = \alpha_i + \beta_1 XAudRating_i + \beta_2 XStg1ExpPic_i + \beta_3 XAudChange_i + \beta_4 XStg1PicChange_i \quad (9)$$

$$XStg2ExpOv_i = \alpha_i + \beta_1 XAudRating_i + \beta_2 XPicRating_i / XStg2ExpPic_i + \beta_3 XAudChange_i + \beta_4 XStg1PicChange_i + \beta_5 XStg2PicChange_i \quad (10)$$

Table 4: Reference-Dependence Variables

Variable	Meaning
$XStg1ExpOv_i/XStg2ExpOv_i$	The expected value for the overall rating of the product as indicated by the participant in stage 1/2.
$XStg1ExpPic_i/XPicRating_i/XStg2ExpPic_i$	In stage 1, none of the participants know the picture quality rating for either product, so $XStg1ExpPic_i$ indicates their expected value for this value. In stage 2, some participants know the value if this was revealed them and some do not. Regressions were run separately in this stage depending on whether X or Y had been revealed to them, and therefore the regression includes $XPicRating_i$ if they were in the self-reveal condition and know this value or $XStg2ExpPic_i$ if they were revealed the value of Y and hence made estimations about the value for X. When these regressions are done for Y, it is $YPicRating_i$ for those in the other-reveal condition and $YStg2ExpPic_i$ if it is the self-reveal.
$XAudChange_i$	This value is equal to the revealed audio quality rating-2.5+a standard normal error term. This term examines how the revealed audio quality rating compares to what had been expected in the previous stage. Since the participant's true expectations were not collected, the revealed rating is compared to the true mean of the audio quality distribution plus a random error term to account for idiosyncratic errors in judgement. This error term is also necessary to eliminate multicollinearity with $XAudRating_i$.
$Stg1PicChange_i$	Similar to $AudChange_i$, this variable is equal to $Stg1ExpPic_i$ -3.291667+a standard normal error term, since 3.291667 was the mean of the distribution given to participants.
$XStg2PicChange_i$	This variable is either equal to $XStg2ExpPic_i - XStg1ExpPic_i$ in the regressions where the picture quality for the product is not known to reflect the change in expectations or it is $XPicRating_i - XStg1ExpPic_i$ and reflects the change from previous expectations to the current knowledge about the product.

If participants exhibit reference-dependent preferences, then the coefficients on $XAudChange_i$, $Stg1PicChange_i$, and $XStg2PicChange_i$ will be positive and significant indicating that when trait quality exceed previous expectations, overall evaluations increase.

c. Loss Aversion

If consumers are not only reference-dependent in their preferences, but also loss averse, then when their expectations are lower than in previous stages, this will have a greater effect on the overall ratings in absolute value than when their expectations are above where they had been in previous stages. Therefore, interaction terms are added to the above regressions such that one coefficient reflects the reference-dependent effect when it is a gain and one reflects how this effect changes if it is a loss.

$$XStg1ExpOv_i = \alpha_i + \beta_1 XAudRating_i + \beta_2 XStg1ExpPic_i + \beta_3 XAudChange_i + \beta_4 XAudLossChange_i + \beta_5 XStg1PicChange_i + \beta_6 XStg1PicLossChange_i \quad (11)$$

$$XStg2ExpOv_i = \alpha_i + \beta_1 XAudRating_i + \beta_2 XPicRating_i / XStg2ExpPic_i + \beta_3 XAudChange_i + \beta_4 XAudLossChange_i + \beta_5 XStg1PicChange_i + \beta_6 XStg1PicLossChange_i + \beta_7 XStg2PicChange_i + \beta_8 XStg2PicLossChange_i \quad (12)$$

Table 5: Loss Aversion Variables

Variable	**Meaning**
$XAudLossChange_i$/ $XStg1PicLossChange$ /$XStg2PicLossChange_i$	Interaction variables between a dummy variable equal to 1 if the change is negative and 0 if it is positive and the change itself.

If the coefficients on $XAudLossChange_i, XStg1PicLossChange,$ and $XStg2PicLossChange_i$ are positive and significant, then that means the effect of this change on expectations is greater when the change is a loss than when it is a gain.

d. Other versus Self Reveal

According to the theoretical model, the consumer is more likely to choose X when the value of Y is revealed than when X is revealed. This holds in all situations except when both X and Y did worse than what was expected of them in the audio quality. The following logistic regression was run separately for each of the 4 groups depending on how X and Y had each performed for the audio quality relative to what had been expected of them based on the prior distributions.

$$Higher_i = \alpha_i + \beta_1 Other_i \quad (13)$$

Table 6: Other versus Self Reveal Variables

Variable	Meaning
$Higher_i$	Dummy variable equal to 1 if $XStg2ExpOv_i > YStg2ExpOv_i$.
$Other_i$	Dummy variable equal to 1 if the participant was revealed the picture quality for Y, be it randomly or strategically.

If the audio quality of both X and Y was worse than was expected, then the coefficient on $Other_i$ should be negative and significant as it is hypothesized that it is more effective to reveal the picture quality for X than for Y in this condition to increase the relative overall rating of X versus Y. For the other three regressions, the coefficient should be positive and significant as it is better to reveal the rating for Y.

e. Strategic Attribute Disclosure Neglect

If participants are explicitly told that the information they are being shown has been chosen in a strategic manner to increase overall evaluations of product X, they should recognize that the picture quality for X is more likely to be revealed if it is good and the picture quality for Y is more likely to be revealed if it is bad. Therefore, if they are shown the value for X, it is more likely that Y is good, and if they are shown the value for Y, it is more likely that X is bad. As such, for those who received the rating of Y, those who were informed strategically should have lower estimates of X compared to those who were informed randomly. Similarly, of those who received the rating for X, those strategically informed should have significantly higher estimates of Y.

To examine this, a series of t-tests were performed to examine a) whether the estimations of X and Y for the strategic reveal group were significantly different from those of the random reveal group and b) whether the true values of the ratings for these groups were different. If the estimations are not significantly different but the true values are, then participants were neglecting to infer information about the quality of a trait based on the fact that the firm strategically did not reveal this value.

6. Results

a. Halo Effect

Participants in the study did appear to exhibit the halo effect. For product X, when the audio rating increases by one standard deviation (1.152 points), this is accompanied by a 0.096 point increase in the picture quality, or by a 0.110 standard deviation increase. However, the consumer's expectations about the picture quality increases by 0.253 points, or 0.377 of a standard deviation. Similarly, for product Y, when the audio rating increases by one standard deviation or 1.131 points, the picture quality increases by 0.144 points or by 0.171 standard deviations, but the estimates increase by 0.271 points or by 0.411 standard deviations. A χ^2 test confirms that the audio ratings predict participants' estimations of the picture quality better than it does the actual rating for both products. For X, the coefficient on the audio quality when regressed on the expected picture quality was different from the coefficient when regressed on the true picture quality rating at the 5% level ($\chi^2 = 3.94, p = 0.0472$). For Y, the coefficients were different at the 10% level ($\chi^2 = 3.11, p = 0.0777$). Since the coefficients were greater on the audio rating in the regressions for the expected picture quality than for the true picture quality, these results support the hypothesis about the halo effect as participants were using their knowledge of the audio quality trait as more informative of the picture quality than is actually true in the data. Therefore, the expectations that participants had about the picture quality following the revelation of the audio quality was inaccurate, which could possibly strengthen the effect of reference-dependent preferences.

Table 7: Halo Effect Regressions

VARIABLES	(1) XStg1ExpPic	(2) XPicRating	(3) YStg1ExpPic	(4) YPicRating
XAudRating	0.219***	0.0832		
	(0.0383)	(0.0533)		
YAudRating			0.240***	0.128**
			(0.0378)	(0.0522)
Constant	2.251***	3.154***	2.273***	3.054***
	(0.109)	(0.151)	(0.108)	(0.150)
Observations	200	200	200	200
R-squared	0.142	0.012	0.169	0.029

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

b. Reference-Dependence

However, the evidence for reference-dependent preferences is much weaker, as few of the terms which compare current and past expectations or knowledge about a products' trait rating are significant. Surprisingly, some of the coefficients are negative, which indicates that exceeding expectations about a given trait's quality actually lowers the overall rating of the product. Nevertheless, the reference-dependent term for the audio quality is significant at the 5% level in both regressions 4 and 5. These terms indicate that when the amount by which the revealed audio quality exceeds prior expectations increases by 1 point, the overall product expectation increases by 1.932 points for the former regression and by 2.047 points for the latter. Therefore, there is a small amount of support for the hypothesis regarding reference-dependent preferences.

Table 8: Reference-Dependence Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	XStg1ExpOv	YStg1ExpOv	XStg2ExpOv	XStg2ExpOv	YStg2ExpOv	YStg2ExpOv
Product	X	Y	X	X	Y	Y
Stage	1	1	2	2	2	2
Is PicRating known?	No	No	Yes	No	Yes	No
XAudRating	1.601**		-0.250	0.412		
	(0.648)		(0.908)	(1.208)		
XStg1ExpPic	5.898***					
	(0.964)					
XAudChange	-0.0528		0.120	1.932**		
	(0.476)		(0.661)	(0.949)		
XStg1PicChange	-0.181		0.487	-0.00183		
	(0.482)		(0.639)	(0.999)		
YAudRating		0.918			-1.791	0.436
		(0.656)			(1.253)	(0.867)
YStg1ExpPic		7.108***				
		(0.914)				
YAudChange		0.456			2.047**	0.846
		(0.469)			(1.017)	(0.596)
YStg1PicChange		-0.872*			-1.574	-0.139
		(0.467)			(1.068)	(0.573)
XPicRating			5.024***			
			(1.575)			
XStg2PicChange			-0.967	-3.268*		
			(1.353)	(1.932)		
XStg2ExpPic				7.352***		
				(1.964)		
YPicRating					8.039***	
					(2.354)	
YStg2PicChange					-4.906***	-1.839
					(1.823)	(1.198)
YStg2ExpPic						6.685***
						(1.404)
Constant	35.86***	33.16***	42.87***	34.49***	36.87***	36.18***
	(2.919)	(2.770)	(4.770)	(6.334)	(7.559)	(4.009)
Observations	200	200	134	66	66	134
R-squared	0.351	0.379	0.201	0.444	0.252	0.358

Standard errors in parentheses. Bold variable name if reference-dependent term.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

c. Loss Aversion

The evidence for loss aversion is slightly stronger. Of the 16 different loss averse interaction, 11 of them are positive, and none of the 5 negative terms are significant. Of the 11 positive terms though, only 2 are significant. These terms correspond to the loss aversion about the audio quality ratings in regressions 2 and 3 in the table 9. For the coefficient in 2 on YAudLossChange, this means that if the audio quality is revealed to be an extra point below (rather than above expectations), the reference-dependent term's effect increases by 2.138 additional points. Similarly, for the coefficient in 3 on XAudLossChange, this means that this negative revelation affects overall ratings in stage 2 by 3.095 additional points compared to a positive revelation of equivalent size. Overall, these results suggest there is some evidence that the change in expectations about the audio quality between the prior stage and stage 1 affect not only the stage 1 overall expectations, but also the second stage expectations. However, these outcomes do lack consistency, and it seems somewhat counterintuitive that there are significant results for a change in expectations between the prior stage and stage 1 present in the stage 2 regressions when there are insignificant results regarding the more recent change in expectations about the picture quality between stages 1 and 2.

Table 9: Loss Aversion Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	XStg1ExpOv	YStg1ExpOv	XStg2ExpOv	XStg2ExpOv	YStg2ExpOv	YStg2ExpOv
Product	X	Y	X	X	Y	Y
Stage	1	1	2	2	2	2
Is PicRating known?	No	No	Yes	No	Yes	No
XAudRating	1.596**		-0.111	0.849		
	(0.647)		(0.926)	(1.252)		
XStg1ExpPic	5.730***					
	(0.980)					
XAudChange	-0.437		-1.289	-0.0169		
	(0.712)		(0.970)	(1.708)		
XAudLossChange	0.787		3.095*	3.305		
	(1.121)		(1.582)	(2.386)		
XStg1PicChange	-1.278		-0.120	2.222		
	(0.928)		(1.296)	(1.790)		
XStg1PicLossChange	1.802		1.295	-4.030		
	(1.246)		(1.738)	(2.513)		
YAudRating		0.712			-2.196	0.364
		(0.657)			(1.404)	(0.874)
YStg1ExpPic		7.066***				
		(0.913)				
YAudChange		-0.427			0.489	-0.0514
		(0.644)			(1.675)	(0.807)
YAudLossChange		2.138*			3.359	2.225
		(1.132)			(2.815)	(1.432)
YStg1PicChange		-1.817*			-0.919	-1.040
		(1.016)			(2.369)	(1.256)
YStg1PicLossChange		1.527			-0.788	1.409
		(1.393)			(3.455)	(1.704)
XPicRating			4.231**			
			(1.654)			
XStg2PicChange			0.137	-3.333		
			(1.472)	(2.570)		
XStg2PicLossChange			-1.293	0.333		
			(2.905)	(3.776)		
XStg2ExpPic				7.779***		
				(1.994)		
YPicRating					9.005***	
					(2.543)	
YStg2PicChange					-5.185**	-1.559
					(2.264)	(1.417)
YStg2PicLossChange					-0.921	-0.455
					(2.987)	(2.165)
YStg2ExpPic						6.298***
						(1.433)
Constant	37.70***	35.79***	47.04***	32.02***	36.06***	39.43***
	(3.213)	(2.975)	(5.392)	(6.520)	(7.853)	(4.472)
Observations	200	200	134	66	66	134
R-squared	0.360	0.397	0.230	0.482	0.271	0.376

Standard errors in parentheses. Bold variable name if loss averse interaction term.
*** p<0.01, ** p<0.05, * p<0.1

d. Other versus Self Reveal

There are no significant results that support the hypotheses regarding how the effectiveness of revealing the picture quality of X versus Y depends on how the audio quality compared to expectations of it. It was hypothesized that when both X and Y had audio quality ratings that were below expectations about them that it would be more effective to reveal the rating for X. While this was true as revealing the value of Y as opposed to X decreased the probability that X would be given an overall rating that was higher in the final stage by 4.5 percentage points, this effect was not significant. For each of the other three scenarios about how the audio quality ratings for X and Y compared to expectations about these values it was hypothesized that it would be more effective to reveal the value of Y. However, the only case where this was true was when X did better than expectations and Y did worse than expectations. In this situation, revealing the picture quality for Y as opposed to X increased the effectiveness by 6.6 percentage points, but again, this effect was insignificant. While this lack of significant results does weaken the model, the hypotheses about the effectiveness of each of the types of revelations were developed based specifically on the significance of the reference-dependence terms about the stage 2 picture quality which correspond to the coefficients on XStg2PicChange, XStg2PicLossChange, YStg2PicChange, and YStg2PicLossChange in regressions 3,4,5, and 6 in table 9. Given the unreliability of those results, the lack of results here should be expected.

Table 10: Other versus Self Reveal Regressions

VARIABLES	(1) Higher	(2) Higher	(3) Higher	(4) Higher
XAud Compared to Prior	Below	Below	Above	Above
YAud Compared to Prior	Below	Above	Below	Above
Other	-0.184	-0.716	0.270	-0.714
	(0.608)	(1.165)	(0.655)	(0.523)
Constant	-0.268	-1.482***	0.241	0.154
	(0.368)	(0.495)	(0.403)	(0.278)
Observations	48	37	41	74

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

e. Strategic Attribute Disclosure Neglect

There is strong evidence for strategic attribute disclosure neglect in the condition where the picture quality for Y was revealed and participants made estimates about the picture quality rating for X. To ensure that the ratings were in fact lower for those in the strategic condition, a t-test was run comparing the true picture quality ratings of X for those who randomly saw the value of Y versus those who strategically saw the value of Y. The picture quality rating was significantly lower for those in the strategic condition at the 1% level ($\mu_{Random} = 3.4000, \mu_{Strageic} = 2.6875, p = 0.0045$) indicating that it would have been rational for those in the strategic group to have lower estimates. In reality, those in the strategic group actually had *higher* estimates about the picture quality in this stage ($\mu_{Random} = 2.9562, \mu_{Strageic} = 3.1263, p = 0.3906$). Though this difference was insignificant, if they had been rational, the strategic group should have had lower estimates, so these results suggest that participants failed to properly respond to the nature by which they are receiving the information in this condition.

The results for those who were revealed the picture quality of X and then made estimates about the picture quality of Y are more ambiguous. Again, there appears to be some irrationality in the estimates since those who received the value of X strategically should theoretically have higher estimates of the value of Y if they are rational since Y is less likely to be revealed the better it is. Yet the estimates of Y were actually slightly higher for those in the random condition ($\mu_{Random} = 3.1715,\ \mu_{Strageic} = 2.9771, p = 0.1403$). While the data about the true picture quality do show that the mean rating for Y was slightly higher for those in the strategic condition, this difference is insignificant, making conclusions on the rationality of this group harder to judge ($\mu_{Random} = 3.4259,\ \mu_{Strageic} = 3.4750, p = 0.7433$). However, the general pattern of these results also support the hypothesis that consumers were neglecting the nature of the information revelation as they were not intuiting the strategy behind the "firm's" decision to display one product's attribute while withholding the other, since they did not properly adjust their estimates about the rating for the unrevealed product.

Table 11: Strategic Attribute Disclosure Neglect T-tests

	MeanRandom	MeanStrategic	p for mu(Ran)-mu(Strat)	N
XStg2ExpPic	2.9562	3.12625	.3906142	66
XPicRating	3.4	2.6875	.0045106	66
YStg2ExpPic	3.171481	2.977125	.140334	134
YPicRating	3.425926	3.475	.7432935	134

7. Limitations

There are three main areas of concern about this experiment. The first one was a design flaw—unfortunately, the specific prior expectations participants had about the audio and picture quality of a randomly chosen camcorder were not collected. Since they were originally presented with the overall distributions, it was expected that participants understood these and therefore would have simply mimicked the given distributions in their estimations. However, this is not necessarily true, and additionally, by just comparing their current estimations of a trait to the mean of the distribution for the reference-dependent term, this created perfect multicollinearity when trying to also include the true value of the trait in the loss aversion regressions. Therefore, a standard normal error term was included to allow for the potential for errors in their estimations and to correct for this multicollinearity. If future work collected their true estimations even after giving participants the distributions, this correction would be unnecessary.

Another area of concern is the true distributions of the audio and picture quality variables. Though products were rated on a 1-5 scaled for each of these traits, by chance, none of the camcorders received a 5 for their audio quality, and none received a 1 for their picture quality. Though the participants were shown these at the beginning of the experiment, they forgot or did not notice it since in their estimations, every participant estimated a non-zero probability that the picture quality was rated a 1 for at least one of their three estimations about picture quality ratings. This is evidence that perhaps future work should include the prior distributions at all stages of the experiment to ensure that if participants are not properly weighting this that it is due to the halo effect bias and not to simply forgetting or not noticing their prior information.

Finally, and most importantly, future research should try and illicit subjective measures of the overall rating. As of now, participants are trying to guess how the raters of the products' overall value may have been affected by their ratings of the individual traits. Therefore, they may be approaching it as a more analytical question and are just trying to be as accurate as possible. However, generally loss aversion is thought of as a subjective, emotional response to negative information. Therefore, if a more subjective measure of participants' overall ratings was elicited, there might be a greater loss averse effect. This might be done by obtaining their willingness to pay for both objects and then randomly giving one participant the object that they had indicated they were willing to pay more for. However, this was impossible in this experiment since Prolific Academic prohibits collecting the personally identifiable information (such as their address) that would have been necessary to do this.

8. Implications

The theory and results from this paper have clear marketing implications. Primarily, firms should pay close care to the initial impressions that consumers have formed about their products. While there are clear benefits when these perceptions are positive, it also comes with drawbacks that need to be considered. Namely, the halo effect means that consumers may have unrealistic expectations about other aspects of the product that could potentially be very damaging to the brand's overall image if they are not met given loss aversion. Therefore, brands may consider devoting even more of their advertising budgets towards negative advertising about their competitors as opposed to revealing or highlighting more aspects of their own product. Finally, the results on the strategic attribute disclosure neglect of participants is good news for advertisers looking to maximize the impact of their campaigns. If consumers did not exhibit this effect, it could attenuate the positive effects of being able to choose what information to highlight. However, the data presented suggest that consumers will not properly understand that they can learn information based on what firms strategically withhold. On behalf of the consumer, these results suggest that measures should be taken to ensure that when judging the quality of traits, these assessments are done in as independent of a manner as possible. Additionally, when firms have the opportunity to strategically disclose information, part of a rational consumer's response to this revelation should include making inferences about the information that was not revealed.

Works Cited

Barton, J., Castillo, M., & Petrie, R. (2016). Negative Campaigning, Fundraising, and Voter Turnout: A Field Experiment. *Journal of Economic Behavior and Organization*, *121*, 99–113. http://dx.doi.org/10.1016/j.jebo.2015.10.007

Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, *78*, 1. https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2

Brown, A., & Fragiadakis, D. (2018). Can Individuals Correctly Respond to Strategic Attribute Disclosure? *Unpublished*. Retrieved from https://drive.google.com/file/d/1IyAQFhv_UE7X4C_IbXwKeNoh2tO7JBlA/view

Cooper, W. H. (1981). Ubiquitous halo. *Psychological Bulletin*, *90*(2), 218–244. https://doi.org/10.1037/0033-2909.90.2.218

Deversi, M., Ispano, A., & Schwardmann, P. (2018). *Spin Doctors: A Model and an Experimental Investigation of Vague Disclosure* (No. 7244). Retrieved from CESifo Group Munich website: https://ideas.repec.org/p/ces/ceswps/_7244.html

Dion, K., Berscheid, E., & Walster, E. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, *24*(3), 285–290. https://doi.org/10.1037/h0033731

Fridkin, K. L., & Kenney, P. J. (2004). Do Negative Messages Work?: The Impact of Negativity on Citizens' Evaluations of Candidates. *American Politics Research*, *32*(5), 570–605. https://doi.org/10.1177/1532673X03260834

Hardie, B. G. S., Johnson, E. J., Peter, & Fader, S. (2007). *Prinred in U S.I MODELING LOSS AVERSION AND REFERENCE DEPENDENCE EFFECTS ON BRAND CHOICE*.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Kalwani, M. U., Yim, C. K., Rinne, H. J., & Sugita, Y. (1990). A Price Expectations Model of Customer Brand Choice. *Journal of Marketing Research*, *27*(3), 251–262. https://doi.org/10.2307/3172584

Koehler, J. J., & Mercer, M. (2009). Selection Neglect in Mutual Fund Advertisements. *Management Science*, *55*(7), 1107–1121. https://doi.org/10.1287/mnsc.1090.1013

Kőszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, *121*(4), 1133–1165.

Landi, M. (2004). Negative campaigning and electoral competition - ProQuest. Retrieved December 18, 2018, from https://search.proquest.com/docview/305144226

Landy, D., & Sigall, H. (1974). Beauty is talent: Task evaluation as a function of the performer's physical attractiveness. *Journal of Personality and Social Psychology*, *29*(3), 299–304. https://doi.org/10.1037/h0036018

Larsen, R. J., & Marx, M. L. (2017). *An Introduction to Mathematical Statistics and Its Applications* (6 edition). Boston: Pearson.

Lau, R. R., Sigelman, L., & Rovner, I. B. (2015). The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment. *The Journal of Politics*. Retrieved from https://www.journals.uchicago.edu/doi/10.1111/j.1468-2508.2007.00618.x

MillerCoors sues Anheuser-Busch over corn syrup ads. (2019, March 22). Retrieved May 1, 2019, from KMGH website: https://www.thedenverchannel.com/money/consumer/millercoors-sues-anheuser-busch-over-misleading-corn-syrup-ads

Rubenstein, L. (2018, October 30). 61% Increase in Volume of Negative Ads. Retrieved December 18, 2018, from Wesleyan Media Project website: http://mediaproject.wesleyan.edu/releases/103018/

Thorndike, E. L. (1920). A constant error in psychological ratings. *Journal of Applied Psychology*, *4*(1), 25–29. https://doi.org/10.1037/h0071663

Verhulst, B., Lodge, M., & Lavine, H. (2010). The Attractiveness Halo: Why Some Candidates are Perceived More Favorably than Others. *Journal of Nonverbal Behavior*, *34*(2), 111–117. https://doi.org/10.1007/s10919-009-0084-z

Versability. (2018, January 23). Top 10 Product Review Websites. Retrieved May 1, 2019, from Thought for Your Penny website: http://thoughtforyourpenny.com/lifestyle/top-10-product-review-websites/

What We Do - Consumer Reports. (n.d.). Retrieved May 1, 2019, from https://www.consumerreports.org/cro/about-us/what-we-do/index.htm

Appendix

a. Summary Statistics

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
XPicRating	200	3.370	0.870	2	5
XAudRating	200	2.595	1.152	1	4
YPicRating	200	3.390	0.843	2	5
YAudRating	200	2.635	1.131	1	4
XStg1ExpPic	200	2.820	0.670	1	4.500
YStg1ExpPic	200	2.904	0.659	1.100	4.530
Other	200	0.330	0.471	0	1
Strategic	200	0.480	0.501	0	1
AudPrior	200	2.419	0.982	0.256	5.069
PicPrior	200	3.322	0.979	-0.0362	5.971
XStg1ExpOv	200	56.73	8.072	38	76
YStg1ExpOv	200	56.68	8.095	38.95	75.55
XStg2ExpOv	200	58.28	8.388	38.60	78
YStg2ExpOv	200	57.66	8.143	38.65	76.85
XStg2ExpPic	66	2.997	0.684	1.260	4.760
YStg2ExpPic	136	3.059	0.743	1.150	4.700
XAudChange	200	0.176	1.496	-4.069	3.620
YAudChange	200	0.216	1.486	-3.208	3.460
XStg1PicChange	200	-0.502	1.263	-3.558	3.866
YStg1PicChange	200	-0.418	1.198	-3.532	2.916
XAudLoss	200	0.425	0.496	0	1
YAudLoss	200	0.445	0.498	0	1
XStg1PicLoss	200	0.675	0.470	0	1
YStg1PicLoss	200	0.610	0.489	0	1
XAudLossChange	200	-0.532	0.838	-4.069	0
YAudLossChange	200	-0.512	0.775	-3.208	0
XStg1PicLossChange	200	-0.789	0.890	-3.558	0
YStg1PicLossChange	200	-0.703	0.877	-3.532	0
XStg2PicChange	200	0.475	0.991	-2.010	3.410
YStg2PicChange	200	0.211	0.925	-2.890	2.900
XStg2PicLossChange	200	-0.176	0.338	-2.010	0
YStg2PicLossChange	200	-0.247	0.478	-2.890	0
Higher	200	0.430	0.496	0	1
XBelowYBelow	200	0.240	0.428	0	1
XBelowYAbove	200	0.185	0.389	0	1
XAboveYBelow	200	0.205	0.405	0	1
XAboveYAbove	200	0.370	0.484	0	1
XPicChange	66	-0.0132	0.673	-2.010	2.490
YPicChange	136	0.232	0.826	-2.890	2.440

b. Payment

In the survey itself, participants were told that they would earn the most money if their estimations were as accurate as possible. Then, the following specific information was included as a footnote that appeared on every page of the survey should they choose to look at it.

Overall Ratings
Each time you make an estimate about the overall rating of a product, you will receive:

10 cents - 0.25*sum for all rating options (((e-t)^2)*pe cents

Where:

- t is the midpoint of the true range that the product was rated within
- e is the midpoint of the range for rating option
- pe is the probability indicated that the rating is in the range

For example:
If the product received a rating of 72, then t=73 (halfway between 71 and 75).

If you indicated there was a 100% chance the product was in the 71-75 range, you would receive:
10-0.25*((73-73)^2*1)=10-0.25*(0^2*1)=10

If you indicated there was a 100% chance the product was in the 76-80 range, you would receive:
10-0.25((78-73)^2*1)=10-0.25(25)= 3.75

If you indicated there was a 33% chance the product was in the 66 to 70 range, 34% chance it was in the 71 to 75 range, and a 33% chance it was in the 76 to 80 range, you would receive
10-0.25*((68-73)^2*.33+(73-73)^2*.34+(78-73)^2*.33)=10-0.25*(8.25+0+8.25)=5.875

Trait Ratings
Each time you make an estimate about the trait rating of a product, you will receive:

10 cents - 4*sum for all rating options (((e-t)^2)*pe cents

Where:

- t is the true rating of the product for that trait
- e is the estimated rating
- pe is the probability indicated that the rating is equal to e

For example:
If the product received a rating of 4 for its picture quality.

If you indicated there was a 100% chance it got a 4, you would receive:
10-4*((4-4)^2*1)=10-4*0=10

If you indicated there was a 100% chance it got a 5, you would receive:
10-4*((5-4)^2*1)=10-4*1=6

If you indicated there was a 10% chance of a 2, 20% chance of a 3, a 40% chance of a 4, and a 30% chance of a 5, you would receive:
10-4*((2-4)^2*0.1+(3-4)^2*0.2+(4-4)^2*0.4+(5-4)^2*0.3)=10-4*(0.4+0.2+0.3)=10-4*0.9=6.4

Notes

-If the formula indicates a negative payment for any of the seven estimations, you will simply receive 0 cents for this decision.
-After all bonuses are taken into account, your earnings will be rounded to the nearest penny.