Learning from Consumer Reviews:
The Role of Selection and Evaluation Biases

Shi (Sherry) Wang
Assistant Professor of Marketing
Fisher College of Business
The Ohio State University
556 Fisher Hall, 2100 Neil Ave
Columbus, Ohio 43210
wang.13362@osu.edu

Anocha Aribarg
Professor of Marketing
Ross School of Business
University of Michigan
701 Tappan Street
Ann Arbor, MI 48109
(734) 763-0599
anocha@umich.edu

Ralf van der Lans
Professor of Marketing
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
+852 2358 7706
rlans@ust.hk

Abstract

The extant marketing literature has demonstrated the impact of consumer reviews on firm performance outcomes at the aggregate level and focused primarily on the role of star ratings, not review content. Only a few papers have investigated how consumers make choice decisions based on review content. In this research, we examine how consumers select a review to read, how they interpret the review's content, and how the review may affect their choices. Inspired by the psychology theory on pre-decision confirmation bias, we conducted an incentive-aligned review-based choice experiment to test i) whether confirmation bias influences consumers to select more positive reviews of a product offered by their preferred vs. non-preferred brand to read (selection bias) and ii) whether they interpret information in favor of their preferred brand's product (evaluation bias). We then incorporate both mechanisms of confirmation bias into a structural choice model with quality learning and show how confirmation bias can manifest as belief updating bias. With the structural model, we further perform counterfactual experiments to show how a retail platform can re-design its review display system to minimize confirmation bias.

In the digital economy era, consumers rely heavily on online reviews to assess different attributes of products to make choice decisions. Some consumer surveys show that 89% of consumers worldwide read reviews before buying a product (Trustpilot 2020), and roughly three quarters of consumers trust consumer reviews as much as they trust personal recommendations (Brightlocal 2022). Extant academic research has also shown that reviews have significant impact on consumer choice and sales performance (see Zheng 2021 for a review). Not surprisingly, some sellers attempt to create fake reviews, and those biased reviews are shown to affect consumer choice (Dellarocas 2006; He, Hollenbeck, and Proserpio 2022; Park, Shin and Xie 2021). Meanwhile, previous psychology literature (Chaxel, Russo and Kerimi 2013; Fraser-Mackenzie and Dror 2009; Russo, Medvec and Meloy 1996) suggests that consumers may be susceptible to pre-decision conformation bias, where they distort information in favor of products that they want to eventually choose. In this paper, we aim to examine whether and how consumers can suffer from pre-decision confirmation bias when they learn about product quality from reviews to make choice decisions. We know that biased reviews lead to biased decisions, but can consumers make biased decisions even when reviews are unbiased?

Chaxel, Russo and Kerimi (2013) propose selection and evaluation biases as the two mechanisms driving pre-decision confirmation bias (shortened as confirmation bias hereafter). How can this bias operate in the context of consumer reviews? Suppose a consumer decides between two products, A and B, and has prior beliefs that favor product A. If the consumer is already inclined to choose product A, she may choose 4 or 5-star reviews for product A and 1 or 2-star reviews for product B to read (i.e., selection bias). What's more, often times review content can be subject to interpretation, particularly for an ambiguous attribute (e.g., ease of use is more ambiguous than photo quality for a camera). This gives the consumer an opportunity to also be biased in interpreting information in a given review in favor of product A (i.e., evaluation bias). Prior consumer review literature has recognized the situation where consumers may discount reviews if they believe that the reviewers may have different taste from them (e.g., Wu et al. 2015; Zhao et. al. 2013). In contrast, we examine if different interpretations can occur even with quality-related attributes due to confirmation bias.

Studying both selection and evaluation biases using observational data is challenging. Despite extant research on consumer reviews, we know little about how consumers select reviews to read. With an exception of Liu, Lee and Srinivasan (2019), previous research cannot track which reviews consumers read, let alone the orders of reading reviews, in observational data. As a result, previous research typically treats reviews as exogenous information and abstracts out review selection. In addition, how reviews are displayed to consumers, filtering algorithms, or star distribution can affect how consumers select the reviews to read. To understand review selection, we need to tease out these confounding contexts. To cleanly study evaluation bias, researchers also need to find evidence that a consumer reading identical content in a review for different products interprets the information differently. Unfortunately, researchers unlikely observe this data pattern in real-world settings.

Given the challenges discussed, we employ an incentive-aligned review-based choice experiment to test i) whether confirmation bias of endowed brand preferences influences consumers to select more positive reviews of a product offered by their preferred vs. non-preferred brand to read and ii) whether they interpret information in favor of their preferred brand's product. The objective of our paper is threefold. First, we document how confirmation bias may affect consumer choice decisions even when reviews are unbiased. Second, we examine the mechanisms underlying the impact of confirmation bias on choice, including review selection and evaluation biases. Finally, because online retailers can potentially manipulate the order of reviews shown to consumers, we conduct counterfactual experiments to show how review display re-design could help minimize confirmation bias.

Our research makes several contributions. First, we propose a conceptual framework to examine the mechanisms in which confirmation bias can distort consumers' use of product reviews to facilitate their quality learning and choice decisions. Second, previous psychology literature on pre-decision confirmation bias has found evidence of evaluation bias but not selection bias. Using a new experimental paradigm, we show that selection bias also exists in a consumer review context. Third, we extend the standard quality learning model by allowing both signals and perceived variance of the signals from consumer reviews to be biased. Such biased signals and perceived signal variance lead consumers to update their beliefs with biased information and give more weights or be more certain about positive (negative) reviews of the preferred (non-preferred) brand's product in updating their beliefs. As a result, we show that confirmation bias can manifest as belief updating bias in a learning model. Finally, we prescribe how online retail platforms can re-design their review display to minimize confirmation bias using counterfactual experiments.

The rest of the paper is organized as follows. First, we review previous literature in relation to our current research. Then we describe our incentive-aligned choice experiment. Next, we present model-free evidence from the experimental data and delignate our structural model. The following sections describe the results from the structural model estimation and counterfactual experiments. We conclude with a discussion of our research contributions, managerial implications, research limitations, and avenues for future research.

Relevant Literature

Our research is at the intersection of the literature on consumer review, confirmation bias, and choice decision with quality learning. In the consumer review literature, researchers have shown that mean ratings from online reviews affect sales (Chevalier and Mayzlin 2006; Dellarocas, Zhang, and Awad 2007), and the effect can be moderated by product and consumer characteristics (Zhu and Zhang 2010) and rating variance (Lee, Bolinger and Staelin 2022; Sun 2012; Wang, Liu and Fang 2015). Most prior consumer review research focuses on ratings, not textual information, and studies the impact of consumer reviews at the aggregate level. Using textual information in movie reviews, Liu (2006) finds only volume and not valence of the reviews to have impact on box office revenues. These conflicting results suggest that consumers may process ratings and review content differently.

Fewer papers have investigated the impact of product reviews on consumer choice. Using regression discontinuity, Vana and Lambrecht (2021) show that individual online reviews influence consumers' purchase likelihood. Using a structural approach, Zhao et al. (2013) propose that consumers learn more from review ratings of a given product than they do from their own past experience of similar products in the context of book purchases. Consumers also update their beliefs about the credibility of reviews based on their own experiences and ratings from reviews on the same books, so that learning from reviews varies across consumers and over time. Using both review ratings and review content, Wu et al. (2015) quantify the value of online reviews for restaurants and find content to be more valuable than ratings. Finally, Liu, Lee and Srinivasan (2019) use machine learning techniques to classify review content and find that aesthetics and price content in consumer reviews have significant influence on purchase conversions across a wide range of product categories. Although some papers have also addressed the questions of how reviews might be manipulated (Dellarocas 2006; Park, Shin and Xie 2021), previous papers assume that consumers unbiasedly process review information. In this paper, we conduct analysis at the individual level and extend previous research by examining how consumers may select different reviews to read and interpret content in the reviews differently because of confirmation bias.

Our paper is inspired by the psychology literature that recognizes consumers' inclination to distort information in favor of products that they have chosen to reduce cognitive dissonance (Elliot and Devine 1994; Cooper and Fazio 1984; Festinger 1962; Jonas et al. 2001). Russo, Medvec and Meloy (1996) and Chaxel, Russo and Kerimi (2013) argue that confirmation bias can occur pre-decision as well. Chaxel et al. (2013) conceptualize that confirmation bias manifests in people's tendency to i) select and ii) evaluate information to reinforce their current beliefs, attitudes, and choices. However, they do not find that consumers search for preference-supporting information before making a choice but find evidence that consumers are positively biased in evaluating preference-supporting information acquired during the search process. The authors used the consumer review context and followed Fraser-Mackenzie and Dror's paradigm (2009) where participants selected reviews to read from a set of 15 review headlines varying from 1 to 5 stars. There are three reviews associated with each number of stars. In our study, we ask participants to select a review of a product with a particular number of stars to read one at a time without information about star distribution nor potential review content (i.e., no headlines shown). We expect that our paradigm will allow us to observe confirmation bias via both selection and evaluation of preference-supporting information.

Our work is also related to the extant literature on choice models with quality learning (Erdem and Keane 1996; see Ching, Erdem and Keane 2013 for a review). These models follow standard Bayesian updating and assume quality signals to be unbiased, and drawn exogenously and independently. It is also conventional to specify all signals to be associated with fixed uncertainty (i.e., signal variance).

Following standard Bayesian updating, Zhao et al. (2013) and Wu et al. (2015) assume that each review is an unbiased quality signal from the population distribution of product quality. However, a few learning models (Kalra, Li and Zhang 2011; Mehta, Chen and Narasimhan 2008) allow quality signals to be biased depending on their sources (e.g., from pharmaceutical companies vs. physicians) and consumers to respond differently to these signals. Mehta, Chen and Narasimhan (2008) aim to capture confirmation bias in their model. However, they rely on a different definition of confirmation bias. In their context, confirmation bias occurs when consumers' experiences with brand quality from their own consumption can be biased by the prior expectation they form from their prior exposure to biased advertising signals. In our research, we do not expect signals from consumer reviews to be biased, but we expect consumers to interpret review content about a product differently depending on their preference for the brand of the product. Using experimental data, we expect to observe consumers interpreting the same review content more positively when they think the review content is associated with the product offered by their preferred brand (i.e., evaluation bias). As a result, although the review signals are unbiased, how consumers encode the signals makes them biased.

Both Zhao et al. (2013) and Wu et al. (2015) construct quality learning models to examine the impact of consumer reviews on choice in the context of experiential products where quality can be confounded with taste preferences (Lee, Bolinger and Staelin 2022). As such, the authors focus on building a model that allows consumers to have different perceived uncertainties associated with different reviews depending on the extent to which reviewers share similar taste to theirs. In our paper, consumers learn about product quality based on two vertically differentiated attributes, photo quality and ease of use. However, we expect consumers to be more certain (i.e., low perceived variance) of positive reviews for the preferred brand's product and of negative reviews for the non-preferred brand's product. Our focus is therefore on capturing differential perceived signal variance driven by confirmation bias, and not mismatched preferences between consumers and reviewers.

In addition, although it is reasonable to expect that consumers may select product reviews to read in a particular manner, Wu et al. (2015) and Liu, Lee and Srinivasan (2019) assume signals from consumer reviews to be exogenous. Other research attempts to address self-selection in reviews but focus on how the review generation, not consumption, can be driven by certain types of consumers adopting a product early or consumers' reviews being influenced by prior critic reviews (Godes and Silva 2012; Ishihara and Liu 2017; Li and Hitt 2008). In our study, we focus on review consumption and expect to observe consumers' tendency to select preference-supporting reviews—positive reviews for preferred brand's product and negative reviews for non-preferred brand's product—to support their choice decisions. Specifically, we allow consumers to endogenously select reviews to read in such a way that they maximize expected utilities in future period.

Our paper is also related to the economic literature that comments on the limitation of Bayesian updating in accommodating confirmation bias. The primary goal of this literature is to understand the phenomenon of belief polarization, where individuals who observe the same information may draw opposite conclusions, and additional information only results in increased polarity. This literature shows that, among other things, belief polarization can be driven by heterogenous prior beliefs (Dixit and Weibull 2007; Acemoglu, Chernozhukov, and Yildiz 2009) and non-Bayesian updating (Rabin and Schrag 1999). Belief polarization is an extreme case, and it is not the goal of our research to study it. Finally, Rabin and Schrag (1999) propose a model of confirmation bias where individuals ignore ambiguous signals that do not conform with their prior beliefs, and thus updating is simply assumed to be biased in the direction of prior beliefs. These authors do not separate selection and evaluation biases. However, we will explore the notion of ambiguous signals by examining how an ambiguous attribute (i.e., ease of use vs. photo quality) may be more subject to evaluation bias.

Incentive-Aligned Choice Experiment

We designed an incentive-aligned review-based choice experiment where we asked participants to imagine that they were considering two specific models of instant camera, each manufactured by a different company (i.e., brand), to purchase from the Amazon website. We provided participants with the description of instant camera and only allowed them to proceed if they were at least "somewhat interested" in using or purchasing an instant camera. Although our focal products, Polaroid Snap and Kodak Printomatic, existed in the market and we obtained their actual product information (no pictures included) and consumer reviews from the Amazon website for our experimental design, we concealed the actual brand and product names in the study. We instead used two neutral brands of Vistaline and Opticon but informed participants that despite the use of fictitious brands, the products involved in the study were real. The use of fictitious brands allows us to cleanly manipulate brand preference (details will follow). The two instant camera models are very similar with respect to their concrete features[1].

Despite their feature similarity, we instructed participants that previous customers had different opinions about these products' photo quality and ease of use, and there were a large number of customer reviews for both camera models on the Amazon website. We informed participants that we had gathered a representative sample of these reviews that focused only on these two aspects for them to read and urged them to use these reviews to help them choose between the two models of instant camera.

To make the review-based choice task incentive compatible, we told participants that in addition to their base compensation, they would also be eligible to enter a raffle where one in every 50 participants could win an instant camera of their choice. If participating, they would receive the camera of their choice should they win the raffle[2]. Every participant agreed to participate.

Prior to exposing them to reviews and asking them to make choices based on the reviews, we manipulated participants' brand preferences by telling them that because we changed their actual brand names to fictitious names, we would provide them with the description of each brand and asked them to use the brand's description in evaluating these brands. We randomly assigned participants to two different brand conditions: Vistaline is preferred to Opticon vs.

---

[1] Both are 10 Megapixel cameras and use 2-by-3-inch ZINK Instant paper; approximately 7.25 ounces in weight and 1 x 4.75 x 3 inches dimension; both cameras are available in 7 colors and have a simple rectangle design; each camera offers a built-in rechargeable Lithium ion battery; both provide a built-in microSD port, allowing you to save photos for later use; both are offered at **the same price.**

[2] Because we used a mixture of reviews from both brands to create a common set of reviews that were randomly shown to participants, we did not want to use their choice from the review task to determine their prize. We instead asked at the end of survey whether they preferred Polaroid Snap or Kodak Printomatic (with pictures of the products) and used their answer to this question to determine their prize.

Opticon is preferred to Vistaline. In the condition where Vistaline is preferred to Opticon, participants read:

"Vistaline and Opticon are well established companies in camera industry. Both of them have been inventing, manufacturing and commercializing instant cameras and related services for a long time. **Vistaline is a pioneer in the industry as an environmentally responsible manufacturer**, suggested by The Corporate Social Ratings Monitor, a 1000+ company industry database. Compared to Opticon, Vistaline is far ahead in adopting Environmental Protection Agency-approved chemicals and procedures in their manufacturing process. Vistaline's factories are equipped with machines that reduce contamination well below threshold values. Committed to protect human health and environment, Vistaline also has a long term cooperation with a non-profit environmental organization to save lands and waters."

For the condition where Opticon is preferred to Vistaline, we simply swapped the two brand names in the brand's description.

Following Russo, Medvec and Meloy (1996), we view brand preference as a source of confirmation bias. That is, we expect consumers to read reviews for products offered by their preferred vs. non-preferred brands differently. The problems of relying on participants' prior preferences for the actual brands are that some participants may have weak preferences for Polaroid and Kodak, and more importantly that their brand preferences may be correlated with product quality. If participants believe that brand signals product quality (ease of use and photo quality in our context) and their evaluation of product quality based on reviews is influenced by their brand preferences, this is rational and cannot be distinguished from confirmation bias. To separately identify confirmation bias, we endow participants with brand preference based on the company's corporate social responsibility, which is supposedly uncorrelated with product quality (Brown and Dacin 1997; Luschs et al. 2010), the approach similar to that adopted by Russo, Medvec and Meloy (1996). We chose this manipulation because company's corporate social responsibility has been shown to affect brand preference (Klien and Dawar 2004)[3].

[3] We thank Itiel Dror for some thoughts about brand manipulation.

After the brand manipulation, we also asked participants to perform a standard conjoint choice task where we showed six triplets of instant camera options, varied by brand (Vistaline vs. Opticon), average photo quality rating based on 1 to 5 stars, average ease of use rating based on 1 to 5 stars, and price ($49, $99 and $149). In manipulating the average ratings, we only vary them at the levels of 3, 3.75 and 4.5, as most average product ratings on the Amazon website is greater than 3. To make the ratings appear more realistic in each conjoint profile, we add a different random generated number between 0 and .5 from a uniform distribution to each of these rating levels for each profile. This conjoint choice task allows us to perform a manipulation check on brand preferences across conditions and measure relative importance of photo quality and ease of use in their decisions to purchase instant cameras across participants.

Before asking participants to read reviews, we measured their prior beliefs about the overall product rating of each instant camera offered by each brand by asking them to rate each camera given what they knew so far where 1 was equivalent to 1 star and 5 was equivalent to 5 star though their rating could have a decimal. Although participants have not learned about photo quality and ease of use of each camera at this point, we expect the brand manipulation to influence the participants' prior beliefs about the two cameras' ratings, as brand preference is a component of product's utility. We will use this measure as another manipulation check.

Next, we asked participants to choose a review to read one at a time. Participants had ten options of reviews to read, each associated with each of the 1 to 5 stars for either Vistaline or Opticon (e.g., Vistaline's camera model – 3 stars). They could read as many reviews as they liked and could also decide to make a choice between the two cameras at any point in time. Note that participants neither knew about the distribution of stars across all reviews nor saw reviews' titles (the design used by Chaxel, Russo and Kerimi 2013; Fraser-Mackenzie and Dror 2009)[4].

We created a set of reviews for participants to read by scraping consumer reviews for Polaroid Snap and Kodak Printomatic from the Amazon website posted in June 2020, filtering for the reviews that mentioned either photo quality or ease of use or both (68% of our reviews discussed both photo quality and ease of use). We then chose ten reviews of different lengths associated with each star. We modified these reviews to get rid of the original brand and product names mentioned in the original reviews, and remove pictures and information that was not related to photo quality or ease of use. The modification was done by accessing the html code of each selected Amazon review page, modifying the review content, and displaying the modified review to be shown to participants in the survey. These modified reviews thus resembled actual reviews consumers would see on Amazon. We ended up with 50 reviews (ten review per star) in total. When participants indicated that they would like to read a review associated with a particular star of a brand, we randomly chose one of ten possible reviews for them to read. Thus, two participants could read an identical review even if they asked to read a review for different brands. This design ensures that our results are not confounded by different review content written for different brands.

After participants read each review, they were asked to:

1) Rate on a scale -50 – 50, "Regardless of the star associated with the review, what is your impression of the product **overall** based on **this particular review you just read**?". This rating captures how participants evaluated the overall quality of a product based on each

[4] Seven participants (1.16%) exhausted all reviews of a star level that we prepared. We included them in the reduced form analysis, but removed them from the structural model analysis.

review. We also used the same scale to ask participants their impression about the product's photo quality and ease of use[5].

2) Rate on a 1-5 scale (a decimal allowed), "After reading all the reviews you have selected so far, how would you rate Brand ___'s camera **overall**?". This rating captures posterior beliefs about the overall rating of the product. We also used the same scale to measure participants' posterior beliefs about the product's photo quality and ease of use. Recall that we had asked participants to rate each camera overall before reading any reviews (i.e., prior belief) as a manipulation check. However, we did not ask this same question at the attribute level.
3) Rate on a 5-point scale, "How certain you are that you receive sufficient information to make your choice decision?

Once participants had made their choice, we asked them to perform another standard conjoint choice task with six triplets of instant camera options, each described by the same attributes and levels as the previous one. However, we showed them a different set of six triplets. We also asked them to rate importance of instant camera attributes, their knowledge about instant camera, whether they had purchased some models of instant cameras including the ones we used in our study, and their demographic information.

Descriptive Statistics and Model-Free Evidence

We collected data from 604 participants via the Lucid Market Place. We set the quota to ensure that the gender and age distribution follows that observed in the general population (see Table 1 for summary statistics). Because we included only participants who were at least somewhat

[5] Participants also chose an option of "The review does not indicate anything related to" either photo quality or ease of use, as some reviews did not discuss both attributes.

interested in instant cameras, they are somewhat knowledgeable on average and about forty percent have purchased an instant camera. We also observe varied numbers of reviews read across participants (median = 3; sd = 4.00; see Figure 1 for the histogram).

Table 1
Summary Statistics

	Mean /Number of Participants	Standard Deviation / Percentage of Participants
Demographics		
Age	40.76	14.92
Gender		
Male	296	49.00%
Female	308	51.00%
Personal Ratings of Importance of Attributes (1=not important at all, 7=very important)		
Brand	4.27	1.85
Ease of Use	5.41	1.34
Battery Life	5.50	1.42
Photo Quality	6.28	1.15
Portability	5.37	1.44
Price	5.37	1.47
Self-Evaluated Knowledge in Instant Cameras (1 = not knowledgeable, 5 = very knowledgeable)		
Knowledge	3.31	0.94
Purchase interest in instant camera		
Very interested	300	49.67%
Interested	169	27.98%
Somewhat interested	135	22.35%
Purchase experience in instant cameras		
Polaroid	105	17.38%
Kodak	63	10.43%
Fuji	146	24.17%
No experience of above three brands	349	57.78%

Figure 1
Number of Reviews Participants Selected to Read

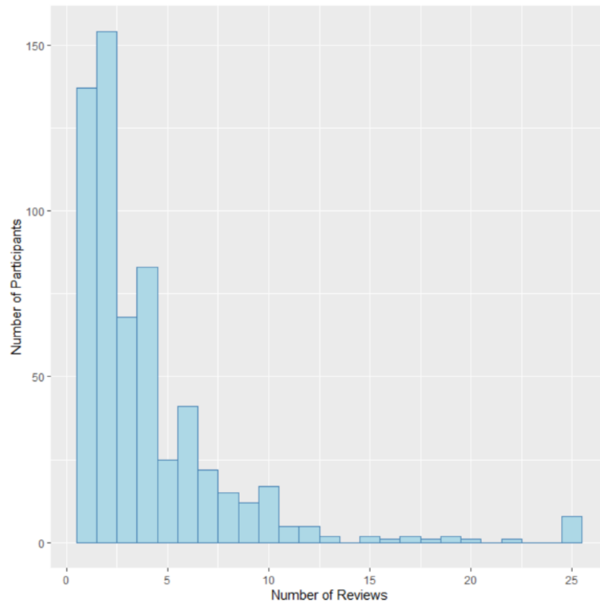

Manipulation Checks

The choice-based conjoint tasks that participants took after the brand manipulation serve as a manipulation check. We refer to the brand that was manipulated to have higher corporate social responsibility as the preferred brand, and the other brand as the non-preferred brand (e.g., Vistaline is the preferred brand and Opticon is the non-preferred brand when Vistaline is more socially responsible). We used the hierarchical Bayes Probit model to estimate each participant's preference for the preferred brand (vs. the non-preferred brand) along with their sensitivities to ease of use, photo quality and price. Table 2 presents the posterior means of all the estimates. On average, participants had a significantly more positive attitude towards the preferred brand. The conjoint estimates also confirm that ease of use and photo quality significantly influenced participants' instant camera choice decisions.

Table 2
Estimates from Choice-Based Conjoint Tasks

Attributes	Brand Preference/Sensitivities
Preferred brand	**.18** (.08, .28)
Ease of use	**.41** (.30, .52)
Photo quality	**1.23** (1.09, 1.38)
Price	**-.30** (-.37, -.22)

Note: 95% posterior intervals are reported in brackets.

Before asking participants to read reviews, we also measured their prior beliefs overall about the instant camera offered by each brand as a manipulation check. We expect that the brand's corporate social responsibility likely influences participants' preferences for the brand and, subsequently, positively affects their ratings for the product offered by the preferred brand. As expected in Figure 2, participants had a significantly higher prior beliefs for the product of their preferred brand than the product of the non-preferred brand (Mean = 4.05 vs. 3.71, p<.001). We also observe that the overall posterior beliefs right before participants made their choice decision are also higher for the preferred brand's product than the non-preferred brand's product (Mean = 3.81 vs. 3.54, p<.001). However, these posterior beliefs are also driven by review selection and evaluation biases. In our structural model, we will separately identify these two components of bias from participants' prior beliefs.

Figure 2
Overall Product Beliefs

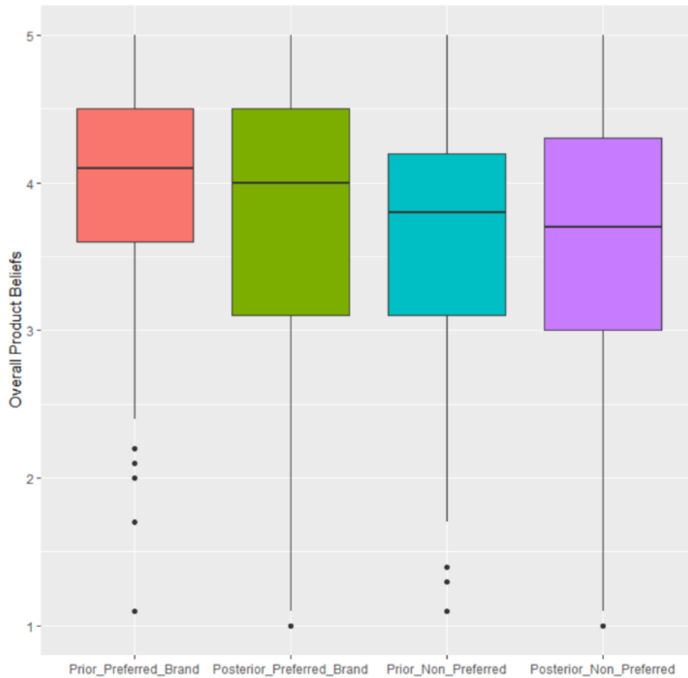

Note: "Prior" represents participants' overall product beliefs after the manipulation and before reading any reviews. "Posterior" represents participants' overall product beliefs after they finish reading reviews and before they make a choice. "Preferred" represents consumers' overall product belief about the product from the preferred brand. "Non" represents participants' overall product belief about the product from the non-preferred brand.

Finally, we observe more positive posterior beliefs about the preferred brand's product than those about the non-preferred brand's product after the last review read. This pattern of posterior beliefs is consistent with the final choice participants made. Approximately 65% (395 out of 604) of participants chose the product from the preferred brand after reading all the reviews they had selected to read. The larger percentage of the preferred brand chosen can be driven by more positive brand preference, as well as confirmation bias. Next, we present model-free evidence for review selection bias and evaluation bias.

Review Selection Bias

We allowed participants to read as many reviews as they would like to help them make their instant camera choice decision in their main review-based choice task. Recall that this choice task is incentive-aligned. On average, participants read more reviews for their preferred brand's product (Means = 2.21 vs. 1.80 reviews, p<.001). Importantly, this difference was mainly driven by participants selecting significantly more positive reviews (4- and 5-star reviews) of their preferred brand's product to read (Table 3).

Table 3
Average Number of Reviews Participants Selected to Read for Each Brand at Each Star Level

	Preferred Brand	Non-Preferred Brand	Difference
1-star	.30	.28	.015 (p=.414)
2-star	.17	.15	.028 (p=.251)
3-star	.34	.28	.060 (p=.060)
4-star	.54	.43	**.118** (p=.031)
5-star	.85	.66	**.192** (p=.003)
Total	2.21	1.80	**.412** (p=.001)

Figure 3
Fitted Loess Line of Chosen Review Brand

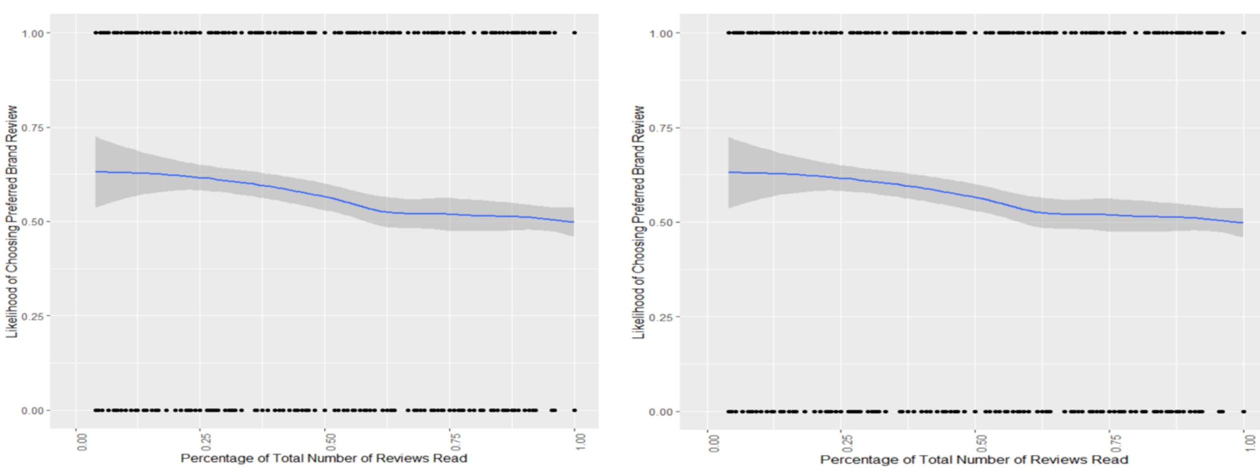

(a) (b)

To explore further, in Figures 3 and 4 we look at the order in which participants selected reviews (preferred vs. non-preferred brand, star levels), given the beliefs they held at a given time. Notably, although participants were manipulated to prefer the product of a socially responsible brand, their beliefs for both products could evolve such that their updated beliefs for the non-preferred brand's product can become more positive. We find that participants were more likely to read reviews from their preferred brand when they were at an earlier stage in their decision process (see Figure 3a) and when their beliefs about the preferred brand's product was more positive than the non-preferred brand's product (see Figure 3b).

Figure 4
Fitted Loess Line of the Star Level of Chosen Reviews

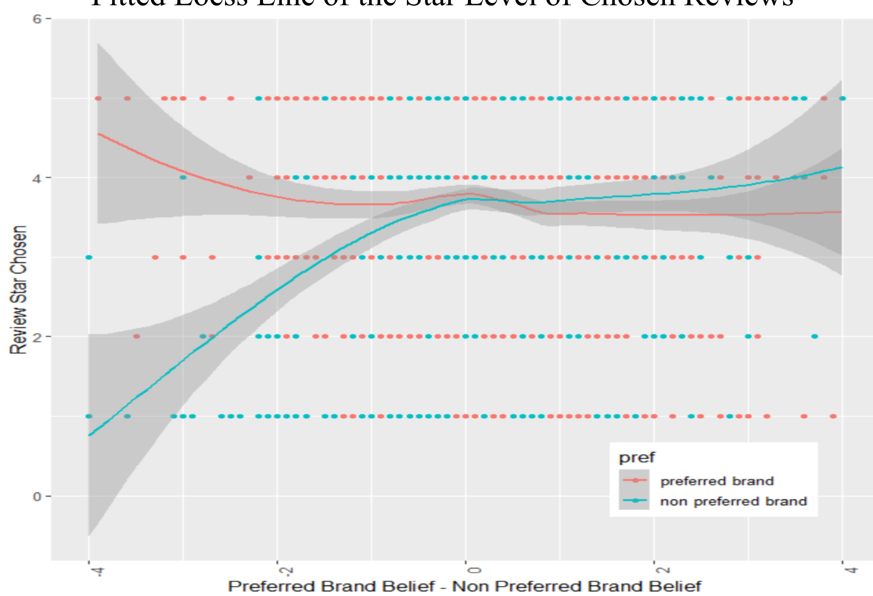

Notes: When the beliefs of two products are close to each other, the value on the x-axis is close to 0. As the preferred brand's product is believed to be better, x becomes more positive and as the non-preferred brand's product is believed to be better, x becomes more negative.

In Figure 4, we examine how participants selected reviews from preferred vs. non-preferred brands to read. On average, participants likely chose positive reviews for the preferred brand to read. Further, when the product of one brand was believed to be worse at a given time, participants tended to choose negative reviews from the better product to read (pink line is above the green line when beliefs of the preferred brand exceed the non-preferred brand, and vice versa). It appears that participants wanted to compensate their beliefs about the product that is perceived to be worse at a particular moment. Interestingly, participants did this to a larger extent when they believed that the preferred brand's product was worse than the non-preferred brand's product (see a larger gap between the green and pink lines towards the left-hand side of Figure 4, as compared to right-hand side). This asymmetric tendency led to unbalanced review selection where participants on average selected more positive reviews for their preferred brand's product and, subsequently, contributed to more positive beliefs participants held about the preferred brand's product.

Review Evaluation Bias

After participants read each review, we asked them to evaluate the impression they got from each review on three aspects: 1) the overall evaluation of the product, 2) ease of use of the product, and 3) photo quality of the product. We find that when reading a review for the product from the preferred brand (vs. non-preferred brand), participants tended to evaluate the review higher in terms of its overall evaluation and its evaluation of ease of use of the product, after controlling for review fixed effects (we included 50 unique reviews in our experiment). However, participants did not evaluate reviews for the preferred brand's product higher in terms of photo quality (Table 4). We conjecture that this pattern of results emerges because photo quality is a more objective attribute than ease of use, and therefore the evaluations for photo quality tend to be less ambiguous. Unlike photo quality, if a review complains about the ease of use the product, participants can attribute it to the reviewer's inability or unfamiliarity with the usage if they want to ignore the review. These results thus suggest that confirmation bias due to review evaluation bias may be more prominent for a more ambiguous attribute.

Table 4
Estimates from Regression of Review Evaluations on Preferred Brand

Variable	Overall Eval	Ease of Use Eval	Photo Quality Eval
Intercept	**-21.625**	**-20.721**	**-24.598**
	(3.233)	**(3.534)**	**(3.573)**
Preferred brand	**2.147**	**2.836**	.318
	(.823)	**(.906)**	(.928)
Review fixed effects	√	√	√
N	2,423	2,325	2,326
R square	.536	.443	.496

Note: Bold denotes significant estimates at 95% confidence interval

Structural Model

To simultaneously characterize the two mechanisms of confirmation bias: selection bias and evaluation bias, we build a structural model in which consumers with different prior beliefs endogenously select reviews to read, interpret the signals they receive from the reviews, and update their beliefs. Our model follows and extends two streams of empirical modeling. The first stream models how consumers learn about product quality/attributes (e.g., match value), and subsequently make choices (Zhao et al. 2013; Wu et al. 2015). The second stream models how consumers endogenously search for information about available products to make optimal choices (Gabaix et al. 2006; Ursu et al. 2022; Yang, Toubia and De Jong 2015).

Consider consumer i who seeks to purchase product $p \in P$. In our experiment, the consideration set P includes two alternatives: p=1 represents the product associated with the preferred brand, and p=2 represents the product associated with the non-preferred brand.

Consumers have uncertainty about the value of two attributes of both products: ease of use and photo quality, which are denoted by k=1 and 2, respectively. We use Q_{pk} to represent the true value of attribute k of product p. To make an informed decision, at each time t consumer i selects to read a review about product p and that has a star rating s, after which the consumer reads a randomly selected review r about product p that has star s. After reading the review, consumer i receives a signal about product attributes ξ_{psr}. She then interprets the signal, updates her belief about product attributes described in review r with respect to product p, and decides whether to read one more review or not. If she decides to read one more review, she will start the process all over again by selecting a review to read. Otherwise, she will make a choice between the two products. Figure 5 presents this consumer decision process. We will describe our model starting from the second step of receiving a signal (i.e., the learning component). Table 5 presents notations we use to denote parameters and variables.

Figure 5
Consumer Decision Process

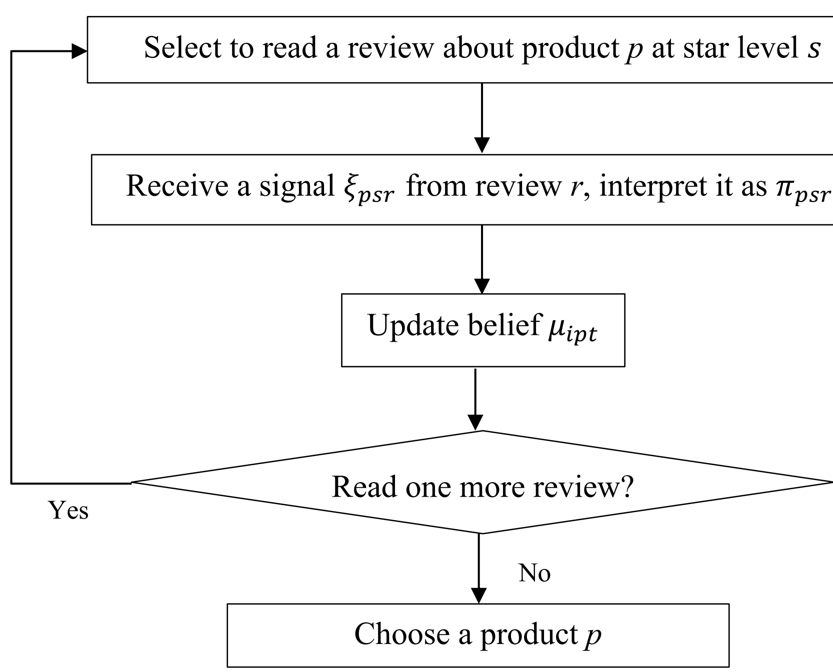

Table 5
Notations of Structural Model

Notation	Meaning
Indexes	
$i = 1, 2, \dots, I$	Consumers
$p = 1, 2, \dots, P$	Products
$k = 1, 2, \dots, K$	Attributes
$t = 1, 2, \dots, T_i$	Decision times
$s = 1, 2, \dots, S$	Star levels of reviews
$r = 1, 2, \dots, R$	Reviews
Learning parameters	
Q_{pk}	True quality of attribute *k* of product *p*
μ_{ipkt}	The mean of consumer *i*'s belief about Q_{pk} at time *t*
b_{ikt}	Consumer *i*'s self-report belief about attribute *k* at time *t*
η_κ^2	Measurement error variance in b_{ikt}
σ_{ipkt}^2	The variance of consumer *i*'s belief about Q_{pk} at time *t*
ξ_{psrk}	The signal that review *r* about product *p* with a star *s* carries about attribute *k*
ρ	Weight of true quality in the mean of the distribution of signals that consumer *i* receives
	Variance of the signal that about review *r*
π_{psrk}	(Biased) interpretation of ξ_{psrk}
e_{ikt}	Consumer *i*'s self-reported evaluation of the review read at time *t* about attribute *k*
η_ι^2	Measurement error variance in e_{ikt}
Selection parameters	
$\tilde{\xi}_{ipskt+1}$	The expected signal that consumer *i* anticipates to receive about product *p*'s attribute *k* at time *t+1* from a *s*-star review
$\tilde{\mu}_{ipskt+1}$	The expected mean of consumer *i*'s belief about Q_{pk} at time *t+1*, given a *s*-star review is read
$\tilde{\sigma}_{ipskt+1}^2$	The expected variance of consumer *i*'s belief about Q_{pk} at time *t+1*, given a *s*-star review is read
$\tilde{u}_{ips,t+1}$	The expected utility of product *p* at time *t*+1 if consumer *s* reads a review from (*p,s)* at time *t*
$\tilde{v}_{ips,t+1}$	The deterministic part of the expected utility of product *p* at time *t*+1 if consumer *i* reads a review from (*p,s)* at time *t*
u_{ipt}	The utility derived by consumer *i* from choosing product *p* at time *t*
v_{ipt}	The deterministic part of the utility derived by consumer *i* from choosing product *p* at time *t*
c_{ipst}	Search cost of consumer *i* to choose to read a review from (*p,s)* at time *t*

Consumer Learning from a Review

When consumers read a review, they can learn about both ease of use and photo quality of the product from reviews. We expect that the magnitude of evaluation bias is influenced by the subjectivity of the attribute. Therefore, we extend the previous literature related to learning (Erdem and Keane 1996; Zhao et al. 2013) from product reviews to allow consumers to learn about two attributes simultaneously. Since the reviews in our experiment are designed to be only about the two attributes, it is natural to assume that consumers' brand preferences (for preferred relative to non-preferred) are only driven by the brand manipulation and do not change as they read reviews about ease of use and photo quality.

At time *t*, consumer *i* reads a review *r* about product p, and has a star rating s. Before time *t*, consumer *i* holds a prior belief about the true quality of attribute *k* of product p, Q_{pk}, which follows a normal distribution $N\left(\mu_{ipk,t-1}, \sigma_{ipk,t-1}^2\right)$. $\mu_{ipk,t-1}$ is the mean, and $\sigma_{ipk,t-1}^2$ captures consumer *i*'s uncertainty about Q_{pk} before time *t*. As mentioned earlier, we chose to manipulate brands in such a way that it should not, or minimally, affect consumers' prior beliefs about ease of use and photo quality for both brands. Given our direct measures of beliefs about the two attributes after each review read, and the updating parameters we backed out from measures of beliefs and signals, we can estimate prior beliefs (before reading any reviews) about both attributes of the two products from our data. Prior research has shown the importance of relaxing the assumption about prior beliefs in search and learning models (Jindal and Aribarg 2021; Ursu et al. 2022); we want to make sure that we tease out the impact of prior beliefs from confirmation bias on choice.

Since each review may contain signals about one or both attributes of the product, we denote the signals carried by this review as $\xi_{psr} = \{\xi_{psrk}\}, k = 1, 2$, where ξ_{psrk} represents the signal from review *r* about product p with a star *s* on attribute *k*. In classic learning models (Ching, Erdem and Keane 2013; Erdem and Keane 1996), the signal consumer *i* receives about product p follows a normal distribution with mean equal to the true quality level Q_{pk}. However, this assumption is not appropriate in our review reading context where consumer *i* knows about the star level of the review before reading it. Intuitively, a five-star review should generate relatively more positive signals and a one-star review should generate relatively more negative signals. Therefore, the signals should be correlated with not only the true quality of the product, but also the star rating of the review. Therefore, we assume that the signal ξ_{psrk} about product p with a star rating *s* follows a normal distribution:

$$\xi_{psrk} \sim N\left((1-\rho)s + \rho Q_{pk}, \tau_{ps}^2\right), \tag{1}$$

where the mean is a weighted average of the true quality Q_{pk} and the star rating s. ρ represents the weight assigned to the true quality level. When ρ approaches 0, the perceived signal mean is purely determined by the star rating, while when ρ approaches 1, the perceived signal mean is purely determined by the true quality. τ_{ps}^2 represents the extent to which each review may deviate from this weighted average.

In classic learning models, the signal variance τ_{ps}^2 is fixed across consumers and over time (Ching, Erdem and Keane 2013; Erdem and Keane 1996), while recent papers allow signal variance to vary (Ursu et al. 2022; Wu et al. 2015; Zhao et al. 2013), that is, it should be viewed as perceived signal variance. For example, Zhao et al (2013) argue that the signal variance increases with the inconsistency between previous purchase experience of books and the reviews of the books provided by other consumers. Ursu et al. (2022) specify that the signal variance associated with a product attribute is affected by a consumer's prior experience with the attribute. To explore confirmation bias in review selection, we allow the perceived signal variance to be influenced by the star rating of the review. Previous research shows empirical evidence that consumers find belief-confirming reviews to be more helpful (Yin, Mitra and Zhang 2016). We hypothesize that the perceived signal variance of review *r* at time *t* depends on the star level s of the review and whether the product p is from preferred or non-preferred brand. As the signal variance is positive, we specify it with an exponential term:

$$log(\tau_{ps}^2) = \theta_{2,s} + \theta_{1,s}(p = 1), \quad (2)$$

where $\theta_{2,s}$ represents the baseline uncertainty of a signal from a s-star review of the non-preferred brand's product. $\theta_{1,s}$ represents the additive uncertainty of a signal from a *s*-star review of the preferred brand's product. A negative value of $\theta_{1,s}$ indicates that a *s*-star review about the preferred brand's product is treated as more informative than a *s*-star review about the non-preferred brand's product. If $\theta_{1,s}$ is negative at higher star levels (i.e., s = 4 or 5), there is selection bias in the sense that a consumer puts higher weight on a positive review about the product from their preferred brand. Because signal variance affects belief uncertainty through Bayesian learning, and belief uncertainty subsequently decreases consumers' product utilities due to risk aversion (Equation 12), consumers will likely select higher star reviews for preferred products to read to reduce uncertainty and maximize their utilities. In our model, we specify review selection bias to emerge because of this biased signal variance.

After receiving a signal following the distribution in Equation (2), consumer *i* updates her belief about the values of the two attributes of the product. Following Kalra, Li and Zhang (2011) and Mehta, Chen and Narasimhan (2008), who allowed signals from certain sources of information to be biased, we assume that consumer i's interpretation of the signal ξ_{psrk} may deviate from the actual signal depending on the product's brand to account for possible evaluation bias. Particularly, the interpretation π_{psrk} she uses to update her belief may deviate from the actual signal ξ_{psrk} by δ_{pk}:

$$\pi_{psrk} = \xi_{psrk} + \delta_{pk}. \tag{3}$$

In our estimation, we fix δ_{2k}, the evaluation bias associated with the non-preferred brand, to be 0 for identification. Recall that in our experiment participants read randomly generated reviews from the same pool regardless of the product. We collected self-reported evaluations of each review consumer *i* reads. Thus, we can observe how consumers may differentially interpret an identical product review, when it is about the product from the preferred brand vs. non-preferred brand. These data are crucial in helping us identify δ_{pk}, which captures the evaluation bias. We further assume that consumer *i* reports their evaluations at *t* with some measurement error $\iota_{ikt} \sim N(0, \eta_\iota^2)$:

$$e_{ikt} = \pi_{psrk} + \iota_{ikt}. \tag{4}$$

Next, we assume that consumers update their beliefs in a Bayesian manner. After receiving a biased signal π_{psrk} with biased signal variance τ_{ps}^2, her belief of Q_{pk} will be updated to have a normal distribution with mean μ_{ipkt} and variance σ_{ipkt}^2, where

$$\mu_{ipkt} = \frac{\frac{1}{\sigma_{ipk,t-1}^2}}{\frac{1}{\tau_{ps}^2}+\frac{1}{\sigma_{ipk,t-1}^2}} \mu_{ipk,t-1} + \frac{\frac{1}{\tau_{ps}^2}}{\frac{1}{\tau_{ps}^2}+\frac{1}{\sigma_{ipk,t-1}^2}} * \pi_{psrk}, \tag{5}$$

$$\sigma_{ipkt}^2 = \frac{1}{\frac{1}{\tau_{ps}^2}+\frac{1}{\sigma_{ipk,t-1}^2}}. \tag{6}$$

Since both biased signal and biased signal variance are used to update consumers' beliefs, consumers' posterior beliefs can be influenced by confirmation bias. Note that unlike most previous learning papers where beliefs are not observed, we also collected self-reported beliefs about ease of use and photo quality after consumers read each review. These belief data help us identify δ_{pk}, $\theta_{1,s}$ and $\theta_{2,s}$. Similar to Equation (4), we assume that consumer *i* reports her mean beliefs at *t* with some measurement error $\kappa_{ikt} \sim N(0, \eta_\kappa^2)$:

$$b_{ikt} = \mu_{ipkt} + \kappa_{ikt}. \tag{7}$$

Consumer Review Selection and Choice

Zhao et al. (2013) and Wu et al. (2015) assume that all reviews associated with a product in the consideration set are read by the consumer. We extend this by explicitly modeling consumers' review selection behavior, and assume that consumers can forward look when deciding whether to read more reviews (i.e., receiving more signals), and which reviews to read. We follow the directed cognition approach to assume that at each period *t*, if the consumer decides to select a review to read, she does so as if she were going to make a choice immediately after reading this review (Gabaix et al. 2006; Yang, Toubia and De Jong 2015). In a context where consumers receive monetary payoff, Gabaix et al. (2006) show that this directed cognition approach predicts better consumers' actual search and decision process, than a fully rational forward-looking model.

Consumer *i* faces uncertainty about Q_{pk}, and she has the opportunity to search for information about both products by selecting to read reviews about these products at different star levels. In our experiment, participants must read at least one review. Therefore, we model that at each time *t* ($t = 2, \ldots, T_i$), consumer *i* can either stop reading and make a choice between the two products, or continue reading one more review. If the consumer decides to continue reading reviews, she can select between two products (*p* = 1, 2) and among five star levels (*s* = 1 to 5). A random review *r* from that category will be displayed to the consumer. We use $y_{it} = (p_{it}, s_{it})$ to represent her review selection decision at time *t*. If the consumer decides to stop reading reviews, then she chooses between the two products, we use $y_{it} = (p_{it}, 0)$ to denote consumer *i*'s final choice of product p_{it} at time *t*.

After time *t*, consumer *i* holds the belief that the true value of attribute *k* of product *p* is distributed $N(\mu_{ipk,t}, \sigma^2_{ipk,t})$. Similar to Equation (1), consumers should expect to receive a more positive signal from a five-star review and a more negative signal from a one-star review. Specifically, from the perspective of the consumer, the expected signal $\tilde{\xi}_{ipsk,t+1}$ delivered by a *s*-star review about product *p*'s attribute *k* should follow a normal distribution with mean being a weighted average of the star rating *s* and true quality level Q_{pk}, which she perceives to have a mean of $\mu_{ipk,t}$, and variance τ^2_{ps}, i.e.,

$$\tilde{\xi}_{ipsk,t+1} \sim N\left((1-\rho)s + \rho\mu_{ipk,t}, \tau^2_{ps}\right), \tag{8}$$

where $log(\tau^2_{ps}) = \theta_{2s} + \theta_{1s}(p = 1)$. We assume consumers to be naïve in the sense that they consider themselves to be objective in evaluating the review and update in a Bayesian manner. The consumer expects that after receiving the expected signal $\tilde{\xi}_{ipsk,t+1}$ from reading a review from (*p, s*), her quality beliefs of product *p*'s attribute *k* would be updated to a normal distribution with expected mean $\tilde{\mu}_{ipsk,t+1}$ and expected variance $\tilde{\sigma}^2_{ipsk,t+1}$, where

$$\tilde{\mu}_{ipsk,t+1} = \frac{\frac{1}{\sigma^2_{ipk,t}}}{\frac{1}{\tau^2_{ps}} + \frac{1}{\sigma^2_{ipk,t}}}\mu_{ipk,t} + \frac{\frac{1}{\tau^2_{ps}}}{\frac{1}{\tau^2_{ps}} + \frac{1}{\sigma^2_{ipk,t}}} * \tilde{\xi}_{ipsk,t+1}, \tag{9}$$

$$\tilde{\sigma}^2_{ipsk,t+1} = \frac{1}{\frac{1}{\tau^2_{ps}} + \frac{1}{\sigma^2_{ipk,t}}}, \tag{10}$$

Utility from product choice. The utility derived by consumer *i* at each time *t* is a function of her current belief and her action. Following Zhao et al. (2013), we assume utility derived by consumer *i* from choosing product *p* at time *t* is:

$$u_{ipt} = v_{ipt} + \varepsilon_{ipt}. \tag{11}$$

The deterministic part is given by:

$$v_{ipt} = \beta_{i0} I(p = 1) + \sum_{k=1}^{2} \left(\beta_{ik} \mu_{ipkt} + \frac{\gamma}{2} \beta_{ik}^{2} \sigma_{ipkt}^{2} \right), \tag{12}$$

where β_{i0} represents consumer i's preference for the preferred brand, which was induced by our manipulation of corporate social responsibility. β_{ik} represents consumer i's sensitivity to attribute k. β_{i0} and β_{ik} are identified from conjoint tasks. Following previous learning models (Erdem and Keane 1996; Zhao et al. 2013), we allow consumers to be risk averse; $\gamma < 0$ represents consumers' risk aversion.

Expected utility driving review selection. Consumer *i* pays a search cost c_{ipst} if she decides to select and read an *s*-star review about product *p* at time *t*. We restrict search cost to be positive, and capture three factors that could affect participants' search decisions: revisiting, comparison, and inertia. First, participants may have an incentive to search for information about a different product or from a different star level than what they have searched already. Therefore, we capture participants' tendency to revisit by incorporating an indicator of whether consumer *i* has selected the same product and the same star level in previous periods, denoted by $Y_{i,t-1} = \{y_{i,1}, y_{i,2}, \dots, y_{i,t-1}\}$. Second, participants may want to compare the two products using reviews at the same star level to be able to evaluate the signal in the review. We capture this comparison mode with an indicator of whether consumer *i* selected the same star of other product at time $t-1$. Lastly, participants' selection of reviews may be state dependent, and we capture this inertia with an indicator of whether consumer *i* selected the same star of the same product at time $t-1$. The search cost of an *s*-star review about product *p* at time *t* is:

$$c_{ipst} = exp\left(\alpha_{c0} + \alpha_{c1} I\left((p,s) \in Y_{i,t-1}\right) + \alpha_{c2} I\left(\left(p \neq y_{i,t-1,1}\right) \wedge \left(s = y_{i,t-1,2}\right)\right) + \alpha_{c3} I\left(\left(p = y_{i,t-1,1}\right) \wedge \left(s = y_{i,t-1,2}\right)\right)\right), \tag{13}$$

where α_{c0} is the baseline search cost. α_{c1} captures participants' reluctance to revisit. A positive value of α_{c1} indicates that participants tend to read reviews from different categories. α_{c2} captures participants' reluctance to compare. A negative value of α_{c2} indicates that participants tend to switch to read a review from the other product but at the same star level to compare. α_{c3} captures participants' inertia. A negative value of α_{c3} indicates that participants tend to read a review from the same product and the same star level.

If consumer *i* decides to stop reading reviews, she chooses the product with the highest perceived utility at time *t*, u_{ipt} (Equation 11). If instead the consumer decides to select another review to read, she will select (*p*, *s*) that maximizes the value of reading one more review before making a choice decision between the two products. Following Yang, Toubia and De Jong (2015) and Ursu et al. (2022), we assume the value of searching comes from being able to select the product that maximizes consumer's expected utility. Specifically, she will choose (*p*, *s*) that maximizes

$$-c_{ipst} + \int_{-\infty}^{\infty} max\left(\tilde{u}_{ips,t+1}, u_{ij,t+1}\right) f\left(\tilde{\xi}_{ips,t+1}\right) d\tilde{\xi}_{ips,t+1} + \varepsilon_{ipst}, \qquad (14)$$

where the first term is the search cost specified in Equation (13). The second term is the consumer's expected utility from choosing the product with the highest utility in the next period *t*+1. Given the realization of signal (for both attributes of product *p*), the consumer may expect to purchase *p* or the other product $j \neq p$, and the expected utility she achieves will be the larger one among $\tilde{u}_{ips,t+1}$ and $u_{ij,t+1}$. Following Equations (11) - (12), $\tilde{u}_{ips,t+1} = \beta_{i0} I(p = 1) + \sum_{k=1}^{2}\left(\beta_{ik}\tilde{\mu}_{ipsk,t+1} + \frac{\gamma}{2}\beta_{ik}^2\tilde{\sigma}_{ipsk,t+1}^2\right) + \varepsilon_{ip0,t+1}$ represents the updated utility from the product *p* whose review is to be read, and $u_{ij,t+1} = \beta_{i0} I(p = 1) + \sum_{k=1}^{2}\left(\beta_{ik}\mu_{ijkt} + \frac{\gamma}{2}\beta_{ik}^2\sigma_{ijkt}^2\right) + \varepsilon_{ij0,t+1}$ is the utility from the other product $j \neq p$.

Choice probability. At time $t = 1$, each consumer selects a review to read. There are in total ten categories she can choose from ($P = 2$ products by $S = 5$ star levels). Assuming ε_{ipst} follows i.i.d. Gumbel distribution (Yang, Toubia and De Jong 2015), the probability of a consumer *i* reading a review $y_{it} = (p, s)$ at time *t* is:

$$Prob(y_{it}) = \frac{exp(V_{ps})}{\sum_{j \in P, l \in S} V_{jl}}, \tag{15}$$

where $V_{ps} = -c_{ipst} + \int_{-\infty}^{\infty} \log\left(\exp(\tilde{v}_{ips,t+1}) + \exp(v_{ij,t+1})\right) f\left(\tilde{\xi}_{ips,t+1}\right) d\tilde{\xi}_{ips,t+1}$ is the expected utility from reading a review (p, s). $f\left(\tilde{\xi}_{ips,t+1}\right)$ is the perceived distribution of the signals delivered by a review from product p with a star rating of s (Equation 8).

At time $t > 1$, each participant can choose to purchase one of the two products or again select a review to read. Therefore, for $t > 1$, there are in total twelve actions she can choose from (one of two products or one of ten review categories). The probability of consumer *i* taking action $y_{it} = (p, s)$ at time $t > 1$ is given in Equation (15), where

$$V_{ps} = \begin{cases} v_{ipt} & \text{if } s = 0 \\ -c_{ipst} + \int_{-\infty}^{\infty} \log\left(\exp(\tilde{v}_{ips,t+1}) + \exp(v_{ij,t+1})\right) f\left(\tilde{\xi}_{ips,t+1}\right) d\tilde{\xi}_{ips,t+1} & \text{if } s \neq 0 \end{cases}. \tag{16}$$

Likelihood function. In sum, the likelihood of observing $\{y_{it}, s_{ikt}, b_{ikt}\}_{it}$ is:

$$L(\{y_{it}, s_{ikt}, b_{ikt}\}_{it}) = \prod_{i,t} Prob(y_{it}, s_{ikt}, b_{ikt}) = \prod_{i,t} \{Prob(y_{it}) \prod_k [Prob(s_{ikt}) \cdot Prob(b_{ikt})]\}, \tag{17}$$

where $Prob(y_{it})$ is specified in Equations (15)-(16), $Prob(s_{ikt}) = f_N\left(s_{ikt}; \pi_{p_{it}s_{it}r_{it}k}, \eta_l^2\right)$, and $Prob(b_{ikt}) = f_N\left(b_{ikt}; \mu_{ip_{it}kt}, \eta_k^2\right)$. $f_N(x; \phi, \varsigma^2)$ denotes the probability density function of a normal distribution with mean ϕ and ς^2 at *x*.

Estimation Results

We first discuss how our data and model specification allow us to identify the parameters, and how we estimate the model. Then we compare three models and present our parameter estimates. Finally, we run counterfactual analysis to study the effect of different review display designs on consumer search and choice.

Model Identification and Estimation

The learning parameters in our model are identified from participants' selection as well as their self-reported evaluations and beliefs. Specifically, recall that after reading each review, we ask participants to report their evaluations of the review and their current beliefs of both attributes of the product. This helps us identify the biased signals π_{psrk} and mean beliefs μ_{ipkt}. The ratio between signal variance τ_{ps}^2 and prior variance σ_{ipkt}^2 is identified from the model specification that participants update their beliefs following Bayes' rule. The weights of prior belief and signal in determining the posterior beliefs help us identify the relative variance of signal and prior belief. Following that, the prior mean belief μ_{ipk0} is also identified from the signal and posterior belief at time *t*=1. The relative evaluation bias $\delta_{1k} - \delta_{2k}$ is identified because we can observe how participants differentially interpret an identical product review, when it is about the product from the preferred brand vs. non-preferred brand. Because we also estimate the objective signals ξ_{psrk}, we fix δ_{2k}, the evaluation bias associated with the non-preferred brand, to be 0 for identification.

We identify the absolute values of signal variance τ_{ps}^2 and prior variance σ_{ipkt}^2 with the additional observation of participants' choice of review to read and our specification of the utility function. Since we only observe one product choice of each participant, we utilize conjoint analysis to identify parameters in the utility function. We use the median individual parameter estimates of brand preference and sensitivities to ease of use and photo quality from the conjoint analysis as β_{ik}. With β_{ik} and participants beliefs μ_{ipkt} identified, the scale of utility is determined. Because belief variance and risk averse parameter appear together in the utility function, they are not separately identified. We choose to estimate the variance parameters τ_{ps}^2 and σ_{ipkt}^2 and normalize the risk parameter γ to -1 (Erdem et al. 2005; Ursu et al. 2022). However, since only the utility difference between the two products is identified, we cannot estimate the prior variance σ_{ipk0}^2 of both products, we therefore specify σ_{ipk0}^2 to be the same across two products.

The search cost c_{ipst} is identified from our observation of when participants stop searching for information, and the observed review reading history helps us identify α_{c1}, α_{c2} and α_{c3}. Parameter ρ, the weight of true quality in the expected signals, is identified from participants' selection of review star to read. When the value of ρ decreases, a five-star review can be expected to lead to a more positive signal and therefore more positive belief in the next period, which give participants incentive to select four- and five-star reviews.

We also conduct simulation study to see if our model is identified. We assign "true" values for all parameters and simulate participants' review selection and product choices. We then estimate the model from the simulated data. The results show that our model is able to recover the "true" parameter values.

Parameter Estimates

We estimated three models, including our proposed model and two nested versions of the proposed model. Model 1 is our proposed model, where we include both selection bias and evaluation bias. Model 2 is our proposed model but without the evaluation bias ($\delta_{1k} = 0$). Model 3 is our proposed model but without the selection bias ($\theta_{1,s} = 0$). Table 6 reports the in-sample fit statistics of the three models. Model 1 performs better than both models 2 and 3, suggesting that incorporating selection bias and evaluation bias improve model performance. Moreover, the improvement of model fit (both AIC and BIC) is larger when incorporating selection bias (the difference between model 1 and 3) than when incorporating evaluation bias (the difference between model 1 and 2). This suggests that in our data, selection bias is more prominent than evaluation bias.

Table 6
Model Fit Comparison

	Model specification	AIC	BIC
Model 1	Proposed model	35,833.77	36,163.03
Model 2	No evaluation bias ($\delta_{1k} = 0$)	35,845.09	36,165.57
Model 3	No selection bias ($\theta_{1,s} = 0$)	35,864.07	36,171.38

We estimate all models using MCMC with Metropolis-Hastings algorithm. Conditional on the parameters in our model, we compute the simulated likelihood of observed data using Equation (17). We draw 12,000 iterations, and use the first 8,000 draws as burn-in. We assured that the algorithm converged by inspecting the plots of the posterior draws. The results of the three models are reported in Table 7.

Table 7
Parameter Estimates

Parameter	Proposed model	No evaluation bias	No selection bias
Prior mean (μ_{pk0})			
Preferred brand's product			
Ease of use	**3.322**	**3.407**	**3.332**
Photo quality	**3.240**	**2.931**	**3.320**
Non-preferred brand's product			
Ease of use	**3.180**	**2.868**	**3.048**
Photo quality	**2.995**	**2.941**	**2.872**
Prior variance (σ_{k0}^2)			
Ease of use	**.461**	**.085**	**.452**
Photo quality	**.265**	**.045**	**.279**
Evaluation bias (δ_{1k})			
Ease of use	**.079**	-	**.127**
Photo quality	-.016	-	.050
Log of Signal variance ($\theta_{p,s}$)			
Intercept (non-preferred brand's product)			
1-star	**-2.844**	**-4.593**	**-2.729**
2-star	**-3.285**	**-5.070**	**-3.317**
3-star	**-3.068**	**-4.809**	**-3.280**
4-star	**-1.619**	**-3.512**	**-1.919**
5-star	**-2.560**	**-4.325**	**-2.882**
Addition of preferred brand's product			
1-star	.031	-.095	-
2-star	-.220	-.311	-
3-star	-.391	-.510	-
4-star	**-.743**	**-.881**	-
5-star	**-.747**	**-.942**	-
Weight of true quality in expected signal (ρ)	**.712**	**.659**	**.723**
Report variances in			
Evaluation (η_ι^2)	**.718**	**.719**	**.718**
Belief (η_κ^2)	**.818**	**.818**	**.820**
Search cost parameters			
Intercept (α_{c0})	**.412**	**.403**	**.413**
Revisit (α_{c1})	**.499**	**.507**	**.499**
Comparison (α_{c2})	**-1.441**	**-1.458**	**-1.434**
Inertia (α_{c3})	**-.856**	**-.851**	**-.857**

Note: Boldface indicates parameter estimates whose 95% posterior interval does not contain 0.

We find that the prior mean beliefs of the preferred brand's product's both attributes ($\mu_{110} = 3.322$ and $\mu_{120} = 3.240$, respectively, for ease of use and photo quality) are slightly higher than the non-preferred brand ($\mu_{210} = 3.180$ and $\mu_{220} = 2.995$, respectively, for ease of use and photo quality), although these differences are also not significant. Therefore, we do not find strong evidence that our corporate social responsibility manipulation affects participants' prior beliefs about the attributes of the preferred vs. non-preferred brand's product.

Importantly, we find support for the two sources of confirmation bias: selection bias and evaluation bias. First, we find that participants have a significant evaluation bias for ease of use ($\delta_{11} = .079$), but not for photo quality ($\delta_{12} = -.016$). This finding echoes our model free evidence and suggests that participants biasedly interpret the signals they receive from reviews, for attributes that are ambiguous and subjective (i.e., ease of use). Second, we find that the signal variance of 4- and 5-star reviews is significantly lower for the preferred brand's product than for the non-preferred brand's product ($\theta_{1,4} = -.743$ and $\theta_{1,5} = -.747$). This suggests that participants attach higher weight (i.e., certainty) to positive reviews of the preferred brand's product, compared to the non-preferred brand's product. This explains that we observe participants choosing more 4- and 5-star reviews to read for their preferred brand's product, which leads to review selection bias.

We also find a significant effect of history on participants' search decisions. In general, people are less likely to revisit. They tend to search a different category of reviews, either a different star rating or the other product, from what they have searched before ($\alpha_{c1} = .499$). Participants are likely to compare. The search cost is relatively lower for reviews about the other product and the same star level in the previous period ($\alpha_{c2} = -1.441$), similar to attribute-based processing of information. Participants exhibit inertia, and the search cost is lower for reviews about the same product and star level as in the previous period ($\alpha_{c3} = -.856$).

Counterfactual Simulations

With the structural model and parameter estimates, we then conduct several counterfactual simulations to explore the impact of different review display designs on consumer information search and choice. We consider four potential designs for a retail platform's review display system. In the first design, consumers freely decide the product and the star of reviews she wants to examine. This is similar to our experimental setting. In real-life scenarios, this is the case if the platform allows consumers to filter for reviews from a specific star level. In the second design, reviews are shown in a random order. This situation is similar to some online platforms' default review display where the most recent review is shown first. In the third design, reviews are shown in an order from 1-star to 5-star for each product. At each decision time, consumers choose which product to search for more information. If it is the first time a consumer examines a product, a 1-star review is shown. If it is the second time a consumer examines a product, a 2-star review is shown. If it is the sixth time a consumer examines a product, we go back to show a 1-star review. This process continues until the consumer makes a choice between the two products. In the fourth design, reviews are shown in an order from 5-star to 1-star for each product.

We use 100 posterior draws of the parameters in our estimation to run simulations for each design. For comparison, for each review display design, we also simulate consumers search and choice decisions if there is no confirmation bias ($\delta_{1k} = 0$, and $\theta_{1,s} = 0$). Table 8 presents the results of the eight counterfactual simulations. In Design 1, when consumers can freely choose the product and the star level they want to read a review from, we find that 60.83% of consumers will choose the preferred brand' product, showing an effect of confirmation bias compared to the "no bias" situation where 56.94% consumers choose the preferred brand's product. Note that we do not expect the choice shares of both products to be equal, as participants were endowed to prefer the preferred brand's product. The average number of reviews consumers read is 3.543 reviews. Compared to Design 1, randomly displaying reviews to consumers (Design 2) decreases the effect of confirmation bias on consumers' choice, with 58.05% consumers choosing the preferred brand's product (compared to 55.12% in the "no bias" situation). Consumers also tend to search for slightly more information than under Design 1 (the average number of reviews read is 3.588 vs. 3.543). The effect of confirmation bias is attenuated even more if the platform displays reviews from 1- to 5-star (Design 3), where 57.35% consumers now choose the preferred brand's product, compared to 56.17% in the "no bias" situation. Consumers on average read 3.725 reviews, more than when the reviews are displayed randomly or by their own choice. On the other hand, displaying reviews from 5- to 1-star (Design 4) amplifies the effect of confirmation bias on consumer choices, with 61.18% consumers choosing the preferred brand's product, compared to 56.17% in the "no bias" situation. Design 4 also reduces consumers' search length with an average consumer reading 3.379 reviews before making a choice.

Table 8
Counterfactual Simulations

	Review display	Bias	% choosing the preferred brand's product	Average number of reviews read	Difference between "Biased" and "No bias"
Design 1	By choice	No bias	56.94%	3.573	3.89%
		Biased	60.83%	3.543	
Design 2	Random	No bias	55.12%	3.587	2.93%
		Biased	58.05%	3.588	
Design 3	From 1 to 5	No bias	56.17%	3.714	1.22%
		Biased	57.35%	3.725	
Design 4	From 5 to 1	No bias	56.17%	3.393	5.01%
		Biased	61.18%	3.379	

Discussion

Previous literature has studied the effect of reviews on sales, with a focus on aggregate data and the effect of review volume and valence (Chevalier and Mayzlin 2006; Liu 2006). There is limited research that studies the effect of reviews on consumer choice (Vana and Lambrecht 2021; Wu et al. 2015; Zhao et al. 2013). Our paper studies the effect of reviews on consumer choice at the individual level. Moreover, our paper considers explicitly that consumers can endogenously seek information from reviews, interpret reviews in a way that is affected by their prior preference, and update their beliefs to make choices. Specifically, we find that consumers positively evaluate ambiguous attributes for their preferred brand's product, and select more positive reviews (4- and 5-star) to read about their preferred brand's product compared to their non-preferred brand's product.

The paper also contributes to the literature of confirmation bias. Previous psychology literature has difficulty claiming the confirmation bias identified from experiments being pre-decisional, because they must ask participants to state their preference or inclination as a seed for confirmation bias. With conjoint tasks, we can check our manipulation without asking consumers to state their preference, which allows us to get a cleaner test of pre-decision confirmation bias. In this context that minimizes the possibility that participants are committed to their preferred product, we find evidence for (pre-decision) confirmation bias.

Furthermore, we identify both selection bias and evaluation bias as the sources of confirmation bias. We attribute our ability to detect selection bias to our experimental where participants chose a review to read at a time without being influenced by review headlines or their perception of the star distribution. We document these biases using model free evidence, and quantify their relative significance of them using a structural model. Methodologically, we extend learning models by allowing signals to be biasedly interpreted, and allowing signals to have different perceived variances in a forward-looking context. Substantively, our counterfactual experiments provide suggestions for how online retail platforms can re-design their review displays to minimize confirmation bias.

While we believe our research makes an important contribution to the consumer reviews, confirmation bias, as well as the learning literature, we acknowledge several limitations of our study. First, in our experiment, participants need to click to examine different reviews, while in a more realistic online setting, people may select to read different reviews by moving their eyes and focus or spend time on certain reviews they are interested in. Future research could run the experiment in a similar setting where we track and model consumers' eye movements as review selection. Second, in our model, we assume that consumers maximize their expected utility from product choice. Under certain conditions, however, consumers may seek information to distinguish the two products. Our model can be extended to assume that consumers select reviews to maximize the expected absolute utility difference between the two products. This will lead to no closed form solution for future expected utility in the forward-looking model. Third, our model assumes that consumers look forward one period, in the sense that if they select a review, they will make a purchase decision in the next period. Despite the computational burden, it might be interesting to study if consumers look forward for more than one period. Fourth, since we observe one selection and choice process for each participant, we are not able to identify learning parameters and evaluation and selection biases at the individual level. It would be interesting to study if evaluation bias and selection biases are heterogenous across people, and who suffer most from confirmation bias.

References

Acemoglu, Doron, Victor Chernozhukov and Muhamet Yildiz (2016), "Fragility of Asymptomatic Agreement under Bayesian Learning," Theoretical Economics, 11(1), 187-225.

Brown, Tom J., and Peter A. Dacin (1997), "The Company and the Product: Corporate Associations and Consumer Product Responses," Journal of Marketing, 61 (January), 68–84.

Brightlocal (2022), "Local Consumer Review Survey 2022", (accessed Oct 31, 2022), [available at https://www.brightlocal.com/research/local-consumer-review-survey/]

Chaxel, Anne-Sophie, J. Edward Russo and Neda Kerimi (2013), "Preference-Driven Biases in Decision Makers' Information Search and Evaluation," Judgment and Decision Making, 8(5), 561-576.

Chevalier, Judith A., and Dina Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews," Journal of Marketing Research, 43(3), 345-354.

Ching, Andrew T., Tülin Erdem, and Michael P. Keane (2013), "Invited Paper—Learning Models: An Assessment of Progress, Challenges, and New Developments," Marketing Science, 32(6), 913-938.

Cooper, Joel, and Russell H. Fazio (1984), "A New Look at Dissonance Theory," Advances in Experimental Social Psychology, Vol. 17., 229-266.

Dellarocas, Chrysanthos (2006), "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms," Management Science, 52(10), 1577-1593.

Dellarocas, Chrysanthos, Xiaoquan Zhang, and Neveen F. Awad (2007), "Exploring the Value of Online Product Reviews in Forecasting Sales: The Case of Motion Pictures," Journal of Interactive Marketing, 21(4), 23-45.

Dixit, Avinash K., and Jörgen W. Weibull (2007), "Political Polarization," Proceedings of the National Academy of Sciences, 104(18), 7351-7356.

Elliot, Andrew J., and Patricia G. Devine (1994), "On the Motivational Nature of Cognitive Dissonance: Dissonance as Psychological Discomfort," Journal of Personality and Social Psychology, 67(3), 382.

Erdem, Tülin, and Michael P. Keane (1996), "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," Marketing Science, 15(1), 1-20.

Erdem, Tülin, Michael P. Keane, T. Sabri Öncü, and Judi Strebel (2005), "Learning about Computers: An Analysis of Information Search and Technology Choice," Quantitative Marketing and Economics, 3(3), 207-247.

Festinger, Leon (1962), "A Theory of Cognitive Dissonance," Vol. 2., Stanford University Press.

Fraser-Mackenzie, Peter AF, and Itiel E. Dror (2009), “Selective Information Sampling: Cognitive Coherence in Evaluation of a Novel Item,” Judgment and Decision Making, 4(4), 307.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg (2006), “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” American Economic Review, 96 (4), 1043-1068.

Godes, David, and José C. Silva (2012), “Sequential and Temporal Dynamics of Online Opinion,” Marketing Science, 31(3), 448-473.

He Sherry, Brett Hollenbeck, and Davide Proserpio (2022), “The Market for Fake Reviews,” Marketing Science, 41(5), 896-921.

Ishihara, Masakazu and Liu, Yuzhou (2017), “A Dynamic Structural Model of Endogenous Consumer Reviews in Durable Goods Markets,” working paper.

Jindal, Pranav, and Anocha Aribarg (2021), “The Importance of Price Beliefs in Consumer Search,” Journal of Marketing Research, 58(2), 321-342.

Jonas, Eva, Stefan Schulz-Hardt, Dieter Frey, and Norman Thelen (2001), “Confirmation Bias in Sequential Information Search after Preliminary Decisions: An Expansion of Dissonance Theoretical Research on Selective Exposure to Information,” Journal of Personality and Social Psychology, 80(4), 557–571.

Kalra, Ajay, Shibo Li, and Wei Zhang (2011), “Understanding Responses to Contradictory Information about Products,” Marketing Science, 30(6), 1098–1114.

Klein, Jill and Niraj Dawar (2004), “Corporate Social Responsibility and Consumers' Attributions and Brand Evaluations in a Product–Harm Crisis,” International Journal of Research in Marketing, 21(3), 203-217.

Lee, Nah, Bryan Bollinger, and Richard Staelin (2022), “Vertical Versus Horizontal Variance in Online Reviews and Their Impact on Demand,” Journal of Marketing Research, *in press*.

Li, Xinxin, and Lorin M. Hitt (2008), “Self-Selection and Information Role of Online Product Reviews,” Information Systems Research, 19(4), 456-474.

Liu, Yong (2006), “Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue,” Journal of Marketing, 70(3), 74-89.

Liu, Xiao, Dokyun Lee, and Kannan Srinivasan (2019), “Large-Scale Cross-Category Analysis of Consumer Review Content on Sales Conversion Leveraging Deep Learning,” Journal of Marketing Research, 56(6), 918-943.

Luchs, Michael G., Rebecca Walker Naylor, Julie R. Irwin, and Rajagopal Raghunathan (2010), “The Sustainability Liability: Potential Negative Effects of Ethicality on Product Preference,” Journal of Marketing, 74 (September), 18–31.

Mehta, Nitin, Xinlei Chen, and Om Narasimhan (2008), "Informing, Transforming, and Persuading: Disentangling the Multiple Effects of Advertising on Brand Choice Decisions," Marketing Science, 27(3), 334-355.

Park, Sungsik, Woochoel Shin, and Jinhong Xie (2021), "The Fateful First Consumer Review," Marketing Science, 40(3), 481-507.

Rabin, Matthew, and Joel L. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias," The Quarterly Journal of Economics, 114(1), 37-82.

Russo, J. Edward, Victoria Husted Medvec, and Margaret G. Meloy (1996), "The Distortion of Information During Decisions," Organizational Behavior and Human Decision Processes, 66(1), 102–111.

Sun, Monic (2012), "How Does the Variance of Product Ratings Matter?," Management Science, 58(4), 696-707.

Trustpilot (2020), "4 Things Every Business Owner Should Know about the State of Reviews", (accessed Nov 1, 2022), [available at https://uk.business.trustpilot.com/reviews/build-trusted-brand/4-things-every-business-owner-should-know-about-the-state-of-reviews?]

Ursu, Raluca, Tulin Erdem, Qingliang Wang, and Zhang, Qianyun Poppy (2022), "Prior Information and Consumer Search: Evidence from Eye-Tracking," working paper.

Vana, Prasad, and Anja Lambrecht (2021), "The Effect of Individual Online Reviews on Purchase Likelihood," Marketing Science, 40(4), 708-730.

Wang, Feng, Xuefeng Liu, and Eric Er Fang (2015), "User Reviews Variance, Critic Reviews Variance, and Product Sales: An Exploration of Customer Breadth and Depth Effects," Journal of Retailing, 91(3), 372-389.

Wu, Chunhua, Hai Che, Tat Y. Chan, and Xianghua Lu (2015), "The Economic Value of Online Reviews," Marketing Science, 34(5), 739–54.

Yang, Liu, Olivier Toubia, and Martijn G. De Jong (2015), "A Bounded Rationality Model of Information Search and Choice in Preference Measurement," Journal of Marketing Research, 52(2), 166-183.

Yin, Dezhi, Sabyasachi Mitra, and Han Zhang (2016), "Research Note—When Do Consumers Value Positive vs. Negative Reviews? An Empirical Investigation of Confirmation Bias in Online Word of Mouth," Information Systems Research, 27(1), 131-144.

Zhao, Yi, Sha Yang, Vishal Narayan, and Ying Zhao (2013), "Modeling Consumer Learning from Online Product Reviews," Marketing Science, 32(1), 153-169.

Zheng, Lili (2021), "The Classification of Online Consumer Reviews: A Systematic Literature Review and Integrative Framework," Journal of Business Research, 135, 226-251.

Zhu, Feng, and Xiaoquan Zhang (2010), “Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics,” Journal of Marketing, 74(2), 133-148.