

Preference-driven biases in decision makers' information search and evaluation

Anne-Sophie Chaxel*

J. Edward Russo[†]

Neda Kerimi[‡]

Abstract

While it is well established that the search for information after a decision is biased toward supporting that decision, the case of preference-supporting search before the decision remains open. Three studies of consumer choices consistently found a complete absence of a pre-choice bias toward searching for preference-supporting information. The absence of this confirming search bias occurred for products that were both hedonic and utilitarian, both expensive and inexpensive, and both high and low in expected brand loyalty. Experiment 3 also verified the presence of the expected post-choice search bias to support the chosen alternative. Therefore the absence of a pre-choice search bias in all three studies was not likely to be due to our using a method that was so insensitive that a search bias would not be observed under any circumstances. In addition to the absence of an effect of prior preferences on information selection, subjects' self-reported search strategies exhibited a clear tendency toward a balance of positive and negative information. Across the three studies, we also tested for the presence of a preference-supporting bias in the evaluation of the information acquired in the search process. This evaluation bias was found both pre- and post-choice.

Keywords: bias, choice, decision making, decision strategies, predecisional distortion, information search, selective exposure.

1 Introduction

Most important decisions involve the use of information that decision makers deliberately search for. A long-claimed shortcoming of such search is a bias toward seeking information that supports the preferred alternative. This confirmation bias in information search is known as “selective exposure” to information (Festinger, 1957). It has traditionally been considered confined to the time after a decision when there is a chosen option that can be supported by additional information.

Hart et al. (2009) performed a meta-analysis of the 91 qualifying studies that tested for the presence and magnitude of this post-choice bias toward selecting confirming information (see also Fischer & Greitemeyer, 2010, for a more theoretical review). They found a mean effect size of $d = .36$, which is equivalent to a ratio of seeking 1.92 pieces of choice-supporting information for every piece of disconfirming information. This meta-analysis summarized and confirmed the presence of selective exposure. It also revealed multiple factors that are associated with more selective exposure (e.g., a strong commit-

ment to an existing belief or the individual characteristic of closed-mindedness) versus less of this bias (e.g., when an existing belief receives support immediately prior to information search or when that information is judged low in quality). Thus, the meta-analysis revealed both the widespread presence of selective exposure and the many moderators that make it a contingent rather than a universal phenomenon. Like many behaviors, its presence and magnitude are sensitive to facilitating or inhibiting characteristics of the environment.

Hart et al.'s (2009) review also presented the main theoretical explanation as the conflict between two goals, defense and accuracy. The goal of defense is achieved by information that supports the chosen alternative (or other maintained attitude, belief, or behavior). It is almost universally accepted as the main driver of selective exposure. The accuracy goal values balanced search in order to achieve a high-quality decision process. It naturally conflicts with the defense goal, which is served by biased search.

For decades, researchers assumed that a confirming search bias was confined to the time after a decision had been made. A chosen alternative was seen as essential because without it there could be no direction toward which to bias any further search. If people have not already chosen an alternative, how can they know which alternative to defend in the search for more information? To complete the argument that selective exposure can only occur post-choice, the countervailing consideration of accu-

The authors thank the Associate Editor, Mike DeKay, for his constructive comments on earlier versions of the manuscript

Copyright: © 2013. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

*McGill University.

[†]Cornell University.

[‡]University of Uppsala.

racy becomes less relevant once the best option has been chosen. More recent work (e.g., Carlson & Guha, 2011; Fischer & Greitemeyer, 2010; Fraser-Mackenzie & Dror, 2009) has begun to investigate whether selective exposure could also be found pre-choice, with theoretical arguments both for and against its presence, as we review shortly.

Note that, while the above distinction between pre-choice versus post-choice is based on time, there is a second distinction that is equally relevant to preference-driven biases. This is the distinction between the two core information processes, the search for relevant information and the subsequent evaluation of the information that is acquired. Although selective exposure is strictly a search phenomenon, biased evaluation may be equally important if for no other reason than its ubiquity. Note that some information is unintentionally encountered. Therefore, while the remaining information is the result of deliberate search, all acquired information is evaluated. Taken together, the two distinctions form a 2 x 2 framework with the two factors of time (pre-choice vs. post-choice) and action (search versus evaluation). Note that biases in the two actions are known, respectively, as selective exposure and information distortion (also called biased assimilation by Lord, Ross & Lepper, 1979). The pre-choice search (selective exposure) cell is the main focus of the present studies, although all four cells are eventually investigated.

1.1 The absence of pre-choice selective exposure

The theoretical argument for why the pre-choice form of selective exposure will not occur relies on changes in the desire for accuracy from pre-choice to post-choice. The claim is that, before a decision has been made, decision makers recognize when they are searching predominantly for confirming information and are aware that such biased search can compromise the quality of their choice. That is, biased search will undermine the accuracy goal of choosing the option with the highest value, and people know this. As a consequence, they will not exhibit selective exposure pre-choice. However, once the decision has been made, subsequent search cannot degrade decision accuracy. Then decision makers are free to search for the confirming information that will satisfy the defense goal. Such confirmation helps them feel better about their choice, including more confidence and fewer second thoughts. Note that at the core of this argument against pre-choice selective exposure is the assumption that the decision makers feel that the decision has not yet been made, so accuracy is still at risk.

1.2 The presence of pre-choice selective exposure

The theoretical argument for the presence of pre-choice biased search requires recognizing the constructive nature of the decision process. As noted above, it has been generally assumed that, before an option is chosen, there is no clear direction toward which to bias search and, hence, no possibility of such a bias. However, decision researchers have come to recognize that a decision is not all-or-none, as appealing as such a dichotomy might be (Keren & Schul, 2009). Instead, preferences are constructed (Payne, Bettman, & Johnson, 1993) and one alternative emerges as the leader in overall preference during the course of the decision process (Brownstein, 2003; Bussemeyer & Townsend, 1993). Indeed, the original version of selective exposure fully accepted its presence prior to a final commitment to one option or position. Historically, selective exposure, both the concept and the term itself, first appeared in the early public opinion literature, e.g., “selective exposure produced by prior attitudes” (Hyman & Sheatsley, 1947, p. 413). In their analysis of voting decisions, Lazarsfeld, Berelson, and Gaudet (1944) stated that “[p]eople selected political material in accord with their own taste and bias. Even those who had not yet made a decision [on their vote] exposed themselves to propaganda, which fit their not-yet-conscious political predispositions” (pp. 79–80). Thus, the original work from which selective exposure in psychology was drawn did accept the possibility of a pre-commitment confirmation bias in information search. Returning to the decision context, as long as decision makers are aware of which alternative is leading pre-choice, they know in which direction to bias the search for more information. This directional knowledge is the essence of the argument for pre-choice selective exposure.

As noted above, the defense goal is almost universally accepted as the main driver of selective exposure, while the accuracy goal would lead to a more balanced search. Based on those motives, Fischer and Greitemeyer (2010) proffered a framework that attempts to account for some inconsistent results, particularly the effect of the accuracy goal to sometimes increase selective exposure. The essence of Fischer and Greitemeyer’s argument was that when the accuracy goal is focused on making a “good” decision, it leads to more “balanced” search, and therefore, to less selective exposure. In contrast, when the accuracy goal is focused not on the decision but on the information (choosing “the qualitatively best pieces of information”), quality is then conflated with consistency between the tentative decision and the new information. This linkage from the accuracy goal to information quality to consistency with the emerging decision leads to more, not less, selective exposure.

1.3 Studies of pre-choice selective exposure

Before reviewing recent studies, we acknowledge three points. First, and paralleling the results for post-choice selective exposure, there is unlikely to be a uniform finding that pre-choice selective exposure either always exists or never exists. There are too many characteristics of a decision environment that can either facilitate or inhibit such a behavior to enable a single result across all task environments. Second, the question of what constitutes a decision is methodologically tricky. Experimenters need to know a subject's currently preferred option in order to specify a direction to test for a confirming bias. At the same time, if subjects' expression of that preference constitutes a decision in their own minds, subsequent search becomes post-choice. Even if we accept that a preference for one option progressively emerges as information is processed and that this preference can direct a search for confirming information that is genuinely pre-choice, how can experimenters discover which option is preferred unobtrusively enough? Said differently, if experimenters instruct subjects to provide a "tentative" decision, is the subsequent search for information to make a "final" decision still pre-choice? The answers to these questions require subjective judgments. Apart from a researcher's choice to be more or less strict in defining what is pre-choice, it is also important to acknowledge the actual differences in how strongly committed decision makers may feel toward a tentative decision. Third, the set of studies reviewed in this section represents only a subset of the existing literature, chosen based on recency of publication and on the studies' estimated relevance to the scope of this paper. Note though that a more complete overview of this stream of research can be found in Hart et al.'s (2009) review of the literature.

Turning now to the recent empirical studies, Fischer et al. (2011) asked subjects to "make a preliminary decision" between two products (e.g., automobiles). Subjects were then offered the opportunity to select information that either confirmed or disconfirmed their preliminary decision. In three studies, the selective exposure bias was observed (though not with statistical reliability in one study).

Fraser-Mackenzie and Dror (2009) presented one product, a cell phone, and user reviews that varied across five "star ratings", where 1-star reviews were very negative and 5-star reviews were very positive. Subjects decided either to buy or not buy the single product offered. In the one condition where the subjects expressed a preference in the form of a "preliminary rating" (based on a list of product specifications), these researchers reported a bias toward selecting confirming user reviews. However, Fraser-Mackenzie and Dror's conclusion is some-

what indirect because they did not attempt to match the star rating of the acquired reviews with each subject's initial preference rating. Instead, the reviews were matched only with the final rating (which was highly correlated with the initial product ratings, $r = .81$). Thus, indirect evidence revealed a significant pre-choice bias toward selective exposure.

Young, Tiedens, Jung, and Tsai (2011) used only subjects who already held the belief that hands-free devices for cellular phones reduce traffic accidents. These subjects were offered a choice of "topic sentences of eight paragraphs, all of which were ostensibly collected from the media", some of which were confirming of accident reduction while others were disconfirming. Subjects selected more confirming (2.00) than disconfirming (1.81) sentences. However, although this difference was in the direction of the selective exposure bias, it was not statistically reliable.

Among the studies that used a tentative decision, the earliest may be Betsch et al. (2001; Experiment 2). However, unlike the three studies cited above, this one used a problem solving task rather than preferential choice. The learning phase of either 15 or 30 trials was designed to indicate an objectively best option. Subjects were then told that the importance of the final decision had increased substantially, which "requires you to consider more detailed information about the quality of the brands before you make your final decision". The available information was a set of market research reports, described by headings that signaled the valence (and some indication of content). Results revealed a significant confirmation bias toward reports that supported the option that subjects had (tentatively) learned was the best one. The bias was larger after 30 learning trials than after 15.

In contrast to the above studies, Carlson and Guha (2011) reported no evidence of pre-choice selective exposure. They discriminated between leader-supporting search and leader-focused search, where the former represented pre-choice selective exposure. Their Study 1 tested for the presence of this pre-choice search bias by allowing subjects to acquire a user review from one of two blogs. The proportion of subjects who chose the blog that supported the leading brand (one of two backpacks) did not differ from the chance base rate.

Why might Carlson and Guha (2011) have observed a null result when other studies found pre-choice selective exposure, at least directionally? At least two possibilities occur. First, might their method of eliciting a tentative preference be sufficiently unobtrusive not to drive what is more like the post-choice version of the search bias? That is, all prior studies directly requested a single, tentative preference, whereas Carlson and Guha used a horse-race metaphor to identify the leading option at multiple points

in the decision process (“Think about this decision as a horse race and the two backpacks as horses in the race. Imagine that the race is already in progress. Which of the backpacks would you consider to be the leader at the current time?”). Although clearly asking subjects to identify the leading option, this less direct horse-race metaphor may have made those subjects feel less like a decision had been made (even if only a “tentative” decision), thus leading to no selective exposure bias. There is a second, very different possible explanation, one resting on the fact that subjects were offered only two sources of information. Fischer, Schulz-Hardt, and Frey (2008) found post-choice selective exposure when 10 items of information were available but not when only 2 were (and there is an equal number of confirming and disconfirming items in both sets, that is, 5–5 or 1–1). Indeed, they found the opposite effect, a bias toward disconfirming items in the 1–1 condition. These researchers claimed that the situation of one confirming and one disconfirming piece of information encouraged decision makers “to be or to appear self-critical” (p. 241). Whether this second, potential explanation for Carlson and Guha’s null finding is correct awaits additional experimentation.

The above review has been confined to search actions, each of which is identifiable as confirming or disconfirming. Thus, we have not included studies that have used eye movements as the information acquisition behavior because both favored and disfavored information receive multiple fixations. The fact that all information has been fixated makes it difficult to conclude that, for instance, more fixations on supporting information also indicate more pre-choice selective exposure (e.g., Glöckner & Herbold, 2011). In a similar fashion, information impact can be inferred from various models like those that assume parallel constraint satisfaction (e.g., Read & Simon, 2012). Some of these studies reveal a significant role for supporting information (e.g., Engel & Glöckner, 2013). Again, however, all of the available information has been acquired, so it is problematic to associate greater importance or impact with a purely search phenomenon like selective exposure.

At least two conclusions might be gathered from these studies. First, the evidence of a pre-choice bias toward searching for confirming information is considerable, but not unanimous. Thus, the potential effect of such environmental factors as the number of pieces of information should not be forgotten. Second, it seems difficult to ignore the possible impact of the strength of a preliminary decision. Its intensity influences decision makers’ sense of whether the subsequent search for information is genuinely pre-choice or is more accurately characterized as post-choice.

1.4 Information search and information evaluation

Investigations of the evaluation of information do not have the long history of the confirmation-driven bias in the search for information (i.e., selective exposure). Nonetheless, they are numerous, especially if they are defined broadly. For instance, post-choice evaluation bias is the second part of Svenson’s differentiation-consolidation theory (Svenson, 1992, 2006; Svenson & Benthorn, 1992; Svenson, Salo & Lindholm, 2009).

There is also considerable work on the pre-choice phase of decision making, particularly on biased evaluation of acquired information. This work is dominated by two closely related streams. The first is based on connectionist theories (e.g., Brownstein, 2003; Read & Simon, 2012; Simon, Pham, Le, & Holyoak, 2001) or constraint satisfaction models (e.g., Glöckner & Betsch, 2008; Holyoak & Simon, 1999; Read & Simon, 2012). The second is the predecisional distortion of information (e.g., Carlson & Pearo, 2004; DeKay, Patiño-Echeverri, & Fischbeck, 2009; Russo, Medvec, & Meloy, 1996; Russo, Carlson, & Meloy, 2006; Wilks, 2002). This work consistently reports the existence of a pre-choice bias in the evaluation of information toward confirming an existing or emerging preference or other belief.

1.5 Plan of the studies

In the following studies, we try to deal with the two concerns expressed in the previous sections. These are selective exposure’s sensitivity to the characteristics of the task environment and the tricky matter of how much commitment to an alternative has already been made when decision makers follow the experimental instructions to reveal their current preference. Regarding the dependence of the results on the task environment, we construct a situation that corresponds to a common real-world environment. Thus, we use mundane consumer choices and draw information from online user reviews. Although this is only one choice environment, at least it is an externally valid one.

Regarding the challenge of limiting the commitment engendered by a preliminary decision, we ask our subjects to report their “leaning” toward one option or the other by moving a slider on a physical scale with no markings other than the endpoints. Our intent is to minimize the expression of a clear intermediate decision. We try to make decision makers’ report of the leading option feel like part of an emerging preference that, quite naturally, cannot be complete until the subsequent information has been acquired and used.

Three studies test for confirmation-driven biases in product choices. The first two studies test for the presence

or absence of preference-driven biases in the pre-choice phase of decisions and also provide some indication of the theoretical explanation of the observed result. The final study examines the entire choice process, both before and after the decision. It attempts (a) to confirm the presence of the familiar post-choice search bias (i.e., selective exposure to information), (b) to replicate the status of a pre-choice search bias observed in the first two studies, and (c) to assess the magnitude of preference-driven biased evaluation (i.e., information distortion) both pre- and post-choice.

2 Experiment 1

Our first study has three goals. First, and fundamental to all that follows, we test for the presence or absence of a confirming search bias prior to the decision. In doing so, we use a relatively unobtrusive tactic for getting subjects to reveal their currently preferred option. Second, we have subjects report the desirability of a tendency to search for confirming or disconfirming information. This addresses one theoretical position, namely that decision makers are aware of the dangers of biased search which, if the decision makers are driven by the accuracy goal, predicts the absence of pre-choice selective exposure. Finally, we check that a confirmation-driven bias in the evaluation of the acquired information is present as usual. That is, we test for the absence of a pre-choice search bias and the presence of a pre-choice evaluation bias.

2.1 Method

Task and materials. All subjects made two decisions, whether to buy or not buy a digital camera and whether to stay or not stay at a resort hotel. These were chosen to be purchase decisions that would normally be consequential (even if hypothetical). They were also selected to represent a more hedonic object (resort hotel) and a more utilitarian one (camera). It is possible that a hedonic product leads to a less analytic choice process. Such a process might be less concerned with balanced information search and, therefore, more susceptible to a confirming search to defend the hedonically preferred option. For each decision, subjects read a cover story in which they were asked to assume (a) that they were in the market for the product, (b) that they had found one that fit their price target, (c) that they were considering purchasing this product, and (d) that they were to search for additional information (actual user reviews) in order to help them make a decision either to buy or not buy it. No brand names, hypothetical or otherwise, were attached to the two products. The camera was identified as Camera

X, while the hotel was referred to as “the hotel”. The only design factor was the counterbalanced order in which the two choices were presented. Because there were no significant differences due to order, this factor is not discussed further.

After the introductory cover story, subjects saw factual information that provided a relatively complete description of each product.¹ The remaining information was available as online user reviews posted in five categories that were identified only by their valence, conveyed as a number of stars. The category headings ranged from 1 star (from a very dissatisfied customer) through 5 stars (from a very satisfied customer). Three reviews were available in each of the five categories, 15 in all (the same stimulus design used by Fraser-Mackenzie & Dror, 2009).


Subjects. From a list of student volunteers at a large North American university, 88 subjects were recruited online. They also completed the study online (but in the laboratory and individually), as they would have if their accessing of user reviews had been for an actual purchase decision. All were compensated for their time with \$5 (n = 86) or course extra credit (n = 2) according to their preferences.

Procedure. After having read the cover story and factual product information, subjects rated their leaning toward buying the product using a continuous slider on a scale with endpoints of “Absolutely sure to NOT buy” and “Absolutely sure to buy”. This “current leaning” scale was subsequently partitioned into 101 points from –50 to 0 (indifference) to 50, though subjects never saw these numerical values. As noted above, this method for recording the current preference was designed to minimize the feeling of commitment to the leading alternative, however officially tentative and reversible that commitment might have been. Following this initial expression of preference, subjects were offered the 15 user reviews. They were required to select and read at least 5 of them, but were invited to read more if they wished. (For 57 of the 176 decisions, subjects searched for more than 5 reviews, but never more than 10.)

After reading a review, subjects provided two responses. First, they rated the review on a “review evaluation” scale from –50 to 50. A rating of the midpoint (0) indicated that the subject judged the review to be neu-

¹For example, the information for the camera was: Price, \$495-\$500; Dimensions, 3.7 x 0.6 x 2.2 in.; Weight, 6 oz.; Resolution, 12 Megapixel; Zoom, Digital 3x, Optical 2x; Memory, 11 MB Internal; Display, LCD display - TFT active matrix - 3.5 in—Color; Max shutter-speed, 1/1000 sec; Misc., Built in flash, Power save, Built-in speaker, Cropping an image, Histogram display, Resizing an image, PictBridge support, Built-in help guide, 16:9 widescreen mode, Touch-screen control, In-camera red-eye fix , USB 2.0 compatibility, 720p HD movie recording, Dynamic Range Optimizer, Blink Detection technology, and Smile Detection technology

Figure 1: Percentages of positive, negative, and neutral information selected by subjects who were leaning against or toward buying the product (Study 1).



tral. Second, they were asked to report on a scale from -50 to 50 how strongly they were leaning toward buying or not buying “considering all the information you have received so far”. These two responses were collected after each selected review. Once subjects had seen their last review, they announced their buy/not-buy decision. Then they expressed their confidence in that decision on a scale from 50 (absolutely not certain) to 100 (absolutely certain).

Subjects then read five “decision strategies” and reported the desirability of engaging in each one. These strategies consisted of two distracters (e.g., “Taking longer than necessary to reach a decision”) followed by the three behaviors of interest to us. These three were searching for confirming user reviews (i.e., selective exposure), searching for disconfirming reviews, and the biased evaluation of the acquired information to support the leading option (also known as the predecisional distortion of information; Russo, Meloy, & Medvec, 1998)². After reading each strategy, subjects were instructed to rate the described behavior’s desirability on a 5-point scale from very undesirable to very desirable. These ratings were intended to assess whether a preference-confirming strategy was recognized as undesirable, both during information search and during evaluation.

2.2 Results

Biased search The distribution of searches, based on prior leaning towards buying, is presented in Figure 1.

²The precise statements were: “People sometimes have a tendency to look mainly for new information that does not challenge their current position or leaning but instead consistently supports it”; “People sometimes have a tendency to look mainly for new information that challenges their current position or leaning because it disagrees with it”; and “People sometimes have a tendency to interpret new information to support their current beliefs. In a decision, this might be ‘seeing’ new information as favoring the direction that you’re leaning toward”.

Subjects were considered to be leaning toward not buying whenever their rating on the current leaning scale was below the midpoint of 0 , and leaning toward buying whenever it was above 0 . Subjects with a leaning of 0 (absence of leaning) are not represented in the figure. Overall, people selected a balance between positive and negative information over their choice process, whether they were leaning toward buying or not.

We performed two tests for the presence or absence of biased search before the decision to buy (or not buy) was made. First, we regressed the valence of the review selected by the subject (number of stars) on the subject’s prior leaning toward buying, on the scale from -50 to 50 (Model 1, see Table 1). For completeness, we also tested in a second regression model (Model 2, see Table 1) for any effects of product type (Hotel coded as -1 , Camera coded as 1) or serial position (i.e., the number of the search, considered as a centered scale variable). Because no predictions were made for any of these additional factors, their inclusion was investigative only. To account for the presence of fixed and random effects, we used a linear mixed model with a Restricted Maximum Likelihood approach (REML) in the statistical package JMP by SAS. More specifically, subject ID and its interaction with product type were specified as random effects in order to control for any correlation within each subject and within each product. Note that we trimmed all searches beyond 5 so that all subjects would make an equal contribution to the analyses. All other observations, except for four missing data points, were included.

In Model 1, the resulting regression coefficient for the impact of prior leaning on the valence of the selected review ($b = -0.003$) was neither positive nor statistically significant ($p = .13$). In Model 2, this estimate remained close to 0 and again did not reach significance ($b = .001$, $p = .48$). Thus, we found no evidence of a bias toward seeking confirming information. Model 2 yielded a significant effect of the type of product, with subjects looking at more positive information for cameras than for hotels ($b = .08$, $p = .01$). However, note that the regression analysis yielded no clear interaction effect between prior leaning and the type of product ($b = -.003$, $p = .08$). Therefore, even if subjects tended to look at more negative information for hotels, this did not result in significantly more biased search. In addition, the variable serial position yielded a significant positive effect on the type of information selected, with positive information sought more as the search progressed. However, this effect was independent of prior leaning, as indicated by the absence of a significant interaction ($p = .30$). Overall, Study 1 found no evidence for a pre-choice bias in the selection of information based on a prior leaning towards buying or not buying.

Table 1: Effect of prior leaning on the selection and evaluation of product information in Study 1 (unstandardized regression coefficients).

| Variables | Model 1 | Model 2 | Model 3 | Model 4 |
|---------------------------------|---------|---------|-----------|-----------|
| Intercept | 2.89*** | 1.88*** | -40.64*** | -41.40*** |
| Prior leaning | -0.003 | 0.001 | .16*** | .14*** |
| Serial position (a) | | .33*** | | 0.4 |
| Product (b) | | .08* | | -3.51*** |
| Valence of the review (c) | | | 13.14*** | 13.04*** |
| Prior leaning x Serial position | | -0.001 | | 0.006 |
| Prior leaning x Product | | -0.003 | | -0.035 |

* $p < .05$, ** $p < .01$, *** $p < .001$; (a) Centered, (b) Camera = 1 and Hotel = -1, (c) Number of stars.

Biased evaluation In contrast to the observed absence of a predecisional confirmation bias in the search for information, the literature presents consistent evidence of a predecisional confirmation bias in the evaluation of the information that is acquired (i.e., predecisional distortion of information). To examine distortion in this research, we used two methods. The first follows the procedure introduced by Meloy and Russo (2004). The second uses regression analyses and enables a more direct comparison of the results of biased evaluation with those of biased search.

Following the method of Meloy and Russo (2004), we first placed unbiased diagnosticity values on the -50-to-50 scale such that each star rating occupied one-fifth of the scale (i.e., equal spacing) and each assumed diagnosticity value was the midpoint of its interval. This yielded the following values: -40 (1-star rating), -20 (2-star rating), 0 (3-star rating, neutral information), 20 (4-star rating), and 40 (5-star rating). Then for each review selected we subtracted its assumed unbiased diagnosticity from the subject's actual rating of that review. If the difference favored the option (buying or not buying) that was leading after the prior piece of information, the absolute difference was signed positively. If the difference favored the opposite option, the absolute difference was signed negatively. Two examples of this computation are given below:

- If the subject was leaning towards buying the product and selected a 5-star review to read that s/he rated 50 on the scale from -50 to 50, then information distortion was calculated to be $|50 - 40|$, signed positively, yielding +10.
- If the subject was leaning towards not buying the product and selected a 2-star review to read that s/he rated as a -24 on the scale from -50 to 50, then information distortion was calculated to be $|-24 - (-20)|$, signed positively, yielding +4.

The average of these scores was then computed for each subject. Across product categories, the mean distortion was 3.92 units.³ For hotels, the grand mean of distortion was 4.52 units on the scale from -50 to 50, which was significantly different from zero ($t(87) = 4.55$, $p < .0001$). For cameras, the grand mean of distortion was 3.34, which also reached significance ($t(87) = 3.53$, $p = .007$). Although the mean distortion for hotels was higher than for cameras (4.52 vs. 3.34), this difference was not statistically reliable ($p = .38$). Overall, these findings show that biased evaluation of information, in contrast to biased search for information, did occur.

To compare the above findings with those obtained for information selection, we ran two parallel regression models (see Table 1). More specifically, the evaluation of information (on the initial scale from -50 to 50) was first regressed on prior leaning (Model 3). This analysis was followed by a more complete regression (Model 4) including not only prior leaning, but also serial position (centered), product type (Camera = 1 and Hotel = -1), and their respective interactions with prior leaning. In addition, because evaluation is strongly correlated with review valence, we included valence as a control variable. As a consequence, any observed effect of leaning on the response variable is independent of the type of review under consideration. Model 3 yielded a strong positive impact of leaning on information evaluation ($b = .16$, $p < .001$). In other words, the more the subjects leaned

³To test for the possible sensitivity of this result to the assumed diagnosticity values, we computed information distortion based on two other plausible sets of diagnostic values. One had more extreme values {-44, -22, 0, 22, and 44} and one less extreme {-30, -15, 0, 15, and 30}, while both preserved equal spacing between adjacent star ratings. For both sets, the computed mean distortion of information was statistically reliable (for both, $p < .0001$). For the first additional set of values, the mean information distortion was 3.62 across product categories; for the second set, $M = 4.08$. Thus, the observed pre-choice bias in the evaluation of information did not depend on the exact diagnosticity values that were assumed.

towards buying, the more the evaluation of information tended to be positive. Model 4 replicated this main effect ($b = .14, p < .001$) and also revealed a main effect of product type in which subjects tended to evaluate information more positively for hotels more than for cameras ($b = -3.51, p < .001$). We could find no explanation for this difference and can only suggest that hotels as a more hedonic product might have led to a desire for more positive information.

Desirability. The theoretical argument for the absence of selective exposure is that individuals are aware that any bias toward searching for confirming over disconfirming information risks compromising the quality of the decision. Subjects rated the desirability of the three decision strategies on a scale from 1 (very undesirable) to 5 (very desirable). These were the tendency to search for confirming information, the tendency to search for disconfirming information, and the tendency to interpret information to support their current position (i.e., predecisional information distortion). The mean response for confirming search was 2.69, a response significantly below the scale midpoint ($t(87) = -3.06, p = .003$) that implied overall undesirability. In contrast, the mean for disconfirming search was 3.52, a value significantly above the midpoint ($t(87) = 5.05, p < .0001$), indicating overall desirability. For completeness, the mean for information distortion was 2.56 ($t(87) = -4.78, p < .0001$). Although subjects did not exhibit a preponderance of either confirming or disconfirming searches, the former was generally considered undesirable while the latter was considered desirable, as confirmed not only by the two independent tests above but also by a matched t test of the two values over the 88 subjects ($t(87) = -5.84, p < .0001$).

Self-reported search rationale. Recall that subjects were asked to report their own search strategies, “Why did you choose the reviews that you read?” Our interest was whether a substantial majority of subjects reported a balanced approach to information search relative to deliberately seeking more positive or more negative user reviews. Two coders who were blind to all hypotheses categorized the responses into four categories: balanced coverage (some positive and some negative reviews), a focus mainly on the positive reviews (i.e., 4-star and 5-star reviews), a focus on mainly negative reviews (i.e., 1-star and 2-star reviews), and a catch-all category of all other responses. Examples of this last category were “specific things from past experiences” and “by intuition”. The two coders agreed on 80% of the responses. All disagreements were resolved by a third coder, who was also blind to the hypotheses. Note that the reported strategies did not permit the identification of preference confirmation or disconfirmation. Thus, our focus was on balanced

search versus unbalanced (i.e., biased) toward either positive or negative reviews. The results confirmed a strong tendency toward balanced search. Of the 84 responses not in the catch-all category, the frequencies of balanced, positive and negative search were, in order, 73, 1, and 10.

2.3 Discussion


The main finding is the complete absence of a pre-choice search bias. There was no evidence of selective exposure to preference-supporting information, a result supported by the reported search strategies. Further, this result is fully compatible with subjects’ average position that seeking confirming information is undesirable (while, interestingly, seeking disconfirming information is desirable). Subjects seem aware of the dangers of a preference-supporting bias in search. We suggest that knowing that confirming search degrades decision accuracy, combined with recognizing when search is confirming versus disconfirming, may explain the absence of selective exposure pre-choice. In contrast, we found a strong bias in the evaluation of information based on prior preferences, thereby replicating previous literature. Note that people are uniformly unaware of this bias in information evaluation (Russo, 2013), whereas people may have a clearer recognition of whether their search is confirming or disconfirming. We suggest that this difference in awareness explains the opposite results for the absence of biased search versus the presence of biased evaluation.

Despite the consistency of the above results, we used only relatively expensive products. In contrast, decision accuracy may matter less for inexpensive products because choosing an inferior product has lower costs. Similarly, products may differ in their ability to induce loyalty or a consistently strong preference for one brand in the category. Loyalty to a brand may induce more support for it in information search and evaluation. To overcome the limited range of products in Experiment 1, we conducted a replication that varied the type of product more systematically.

3 Experiment 2

This follow-up to our first study is intended to test the finding of no selective exposure in a more complete range of products. First, we include inexpensive as well as expensive products. Although the likelihood of an inferior decision may be the same for both price levels, the cost of a mistake is much lower when the product is inexpensive. Thus, we ask whether the same absence of a pre-choice search bias would occur for low-price categories. At the same time, we varied the expectation of brand loyalty, choosing some product categories that typically ex-

Figure 2: Percentages of positive, negative, and neutral information selected by subjects who were leaning against or toward buying the product (Study 2).



hibit high loyalty and others of low brand loyalty. The former may lead subjects to search for more preference-supporting information instead of a balanced information search. These two product characteristics yield an array of four types of products. Thus, the main goal of Experiment 2 is to test for the absence of a pre-choice search bias across a set of products that are both expensive and inexpensive and both high and low in brand loyalty.

3.1 Method

Task and materials Four product categories were tested: a laptop represented the expensive, high-loyalty category; a refrigerator represented the expensive, low-loyalty category; a soft drink represented the inexpensive, high-loyalty category; and a light bulb represented the inexpensive, low-loyalty category. As before, subjects read a cover story for each product, saw several pieces of factual product information, and had to access at least five user reviews identified only by their star rating. Three reviews were available in each of the five star categories, 15 in all. The only design variable was the counterbalanced order in which the four product decisions were presented. Again, there were no significant differences due to order, so this factor is not discussed further.

Subjects From a list of student volunteers at a large North American university 36 subjects were recruited online. They also completed the study online. All were compensated for their time with \$5 ($n = 24$) or course extra credit ($n = 12$) according to their preferences.

Procedure The procedure was identical to that of the first study. The only change was the product categories tested.

3.2 Results

Biased search. Similar to Experiment 1, the distribution of searches, based on prior leaning towards buying, is presented in Figure 2. Overall, people selected a balance between positive and negative information over their choice process, whether they were leaning toward buying or not.

We repeated Experiment 1's mixed linear model of the effect of the prior leaning on the valence of the information selected (number of stars) in two models (see Table 2). Model 1 includes only prior leaning as a predictor, while Model 2 is a more complete model with serial position (centered), expense ("low" coded as -1, "high" coded as 1), loyalty ("low" coded as -1, "high" coded as 1), and relevant interactions. Note that because one subject did not complete the "Fridge" questionnaire, five data points are missing.

We found no positive effect of prior leaning on information search in Model 1 ($b = -.004, p = .09$) or in Model 2 ($b = -.0006, p = .79$), thereby replicating the results of Study 1. In Model 2, the slope of serial position on information selection was positive and significant ($b = .30, p < .001$), so that subjects increasingly looked for positive information over time. Repeating the finding in Study 1, the interaction between serial position and leaning remained non-significant, indicating that the absence of a confirmation bias in information selection holds throughout the choice task. Thus, fully replicating our first study, there was no evidence of a confirmation-driven bias toward searching for information that supported the current leaning.

Biased evaluation In contrast to biased search, the computation of information distortion yielded a significant value (mean = 3.72, $t(35) = 4.24, p < .0001$).⁴ This level of predecisional distortion was similar to the 3.92 observed in Study 1. We repeated the same two REML models with the evaluation of information as a response variable, and controlling for the valence of the review. Model 3 yielded the expected main effect of prior leaning ($b = .13, p < .001$), so that the more subjects were leaning towards buying, the more they evaluated incoming information positively, thereby confirming the presence of a pre-choice bias in information evaluation. Model 4 yielded the same effect ($b = .13, p < .001$). Note that we did not find any effect related to product price or to prod-

⁴We also tested the same two alternative sets of assumed unbiased diagnostic values as in Experiment 1, namely $\{-44, -22, 0, 22, \text{ and } 44\}$ and $\{-30, -15, 0, 15, \text{ and } 30\}$. For both sets, the computed mean distortion of information was statistically reliable (for both, $p < .0001$). For the first additional set of values, the mean information distortion was 3.86 across product categories; for the second set, $M = 3.38$. Again, the observed pre-choice bias in the evaluation of information did not depend on the exact diagnosticity values that were assumed.

Table 2: Effect of prior leaning on the selection and evaluation of product information in Study 2 (unstandardized regression coefficients).

| Variables | Model 1 | Model 2 | Model 3 | Model 4 |
|---------------------------------|---------|---------|-----------|-----------|
| Intercept | 3.02 | 2.12*** | -40.72*** | -41.60*** |
| Prior leaning | -0.004 | -0.0006 | .13*** | .13*** |
| Serial position (a) | | .30*** | | 0.2 |
| Expense (b) | | 0.01 | | -1.11 |
| Loyalty (c) | | -0.002 | | 0.7 |
| Valence of the review (d) | | | 12.83*** | 12.80*** |
| Expense x Loyalty | | 0.005 | | 0.89 |
| Prior leaning x Serial position | | -0.001 | | -.04* |
| Prior leaning x Expense | | 0.004 | | -0.04 |
| Prior leaning x Loyalty | | -0.003 | | 0.01 |

* $p < .05$, ** $p < .01$, *** $p < .001$; (a) Centered, (b) High = 1 and Low = -1, (c) Number of stars.

uct loyalty. An unexpected interaction appeared between prior leaning and serial position ($b = -.04$, $p = .03$), so that the impact of leaning on information evaluation tended to decrease over time. We do not have an explanation for this relation.

Desirability Two subjects did not complete the desirability ratings and are therefore not included in this analysis. The reason for the absence of a confirmation bias in information search seemed to be the same as in Experiment 1, at least as far as our data allowed a conclusion. Subjects' ratings of the desirability of the three "decision strategies" exhibited a similar pattern of means. The two confirming biases were both rated as undesirable on average (for search, $M = 2.74$; for evaluation, $M = 2.56$). In contrast, the mean desirability rating of disconfirming search was 3.32. Note that, in contrast with Study 1, only the desirability of information distortion was significantly lower than the scale mid-point ($t(33) = -2.33$, $p = .026$), while the other two scores were not different from this neutral value ($p = .10$ for disconfirmatory search and $p = .17$ for confirmatory search). However, again a matched t test between desirability of confirming search and disconfirming search yielded the expected significant difference ($t(33) = -2.14$, $p = .039$). Thus, although the desirability ratings of biased search and evaluation did not reach statistical significance, taken together, they were reliably different from each other.

Search Strategies. The self-reported search strategies were coded as in Experiment 1 by the same two coders (proportion of agreement = .91) with disagreements resolved by the same third coder, again all three blind to the hypotheses. From the 36 subjects, 33 responses were

codable, 4 of which fell into the catch-all category. Of the remaining 29 subjects, 25 reported a search strategy that was balanced, 4 a strategy that favored negative reviews, and none a strategy that favored positive reviews.

3.3 Discussion

Experiment 2 generalized the finding from Experiment 1 of no pre-choice selective exposure to products that are inexpensive or vary in expected brand loyalty. There was again no evidence of a preference-supporting search bias prior to the decision to buy or not to buy a product.

However, in both of our studies the valence of the information provided was explicit and precise. There is little ambiguity in 1 to 5 stars as the labels for information categories. Might subjects show more pre-choice selective exposure if the information were less obviously valenced? Might they take advantage of even limited ambiguity to indulge in a search for preference-supporting information? In addition, we have not verified that the usual post-choice confirming search bias is observed, even though the pre-choice version of this bias does not occur. Such a verification of a well-known finding should add credibility to our method for investigating search. In addition, introducing a post-choice phase allows a complete view of decision makers' use of information, both search and evaluation and both pre-choice and post-choice.

4 Experiment 3

Our last study has three goals. First, we ask whether the finding of no pre-choice confirming search bias would oc-

cur when the product information is drawn from online sources that are less clearly valenced than 1 to 5 stars. To this end, we use manufacturer and third-party online sources where the expected valence is either supportive or denigrating of the product, but where these sources might still contain some posts with a neutral or opposite valence.

Our second goal is to verify that a confirmation search bias exits after the product choice even though one does not occur before choice. That is, we sought to eliminate the possibility that, for some unidentified reason, our subjects would never bias their search for information toward support of the current preference, neither before nor after the decision had been made. This second goal requires a post-choice phase and enables a more complete picture of the use of information in a decision.


The third goal is to assess both the pre- and post-choice evaluation of the acquired information. Of particular interest is whether any bias in the post-choice evaluation of information exceeds its pre-choice counterpart, a relation that would parallel what is expected for information search. Alternatively, post-choice distortion may also be lower than pre-choice distortion, a finding that would mirror Russo et al.'s prior findings (1998). This also completes the 2x2 array of information search and evaluation, both pre- and post-choice.

4.1 Method

Task and materials Only one product was used, the laptop computer from our second study. The stimuli consisted of four categories of product information, two likely to be positively valenced and two likely to be negatively valenced. The positive sources were the manufacturer's brand website and an independent fan website for the product. The two negative sources were a blog opposed to Brand X and a retailer's website (Amazon.com) where customers could post their product experiences. The fan page and the company's website were described as likely to contain positive reviews. The blog and the online retailer were described as likely to contain negative reviews. In addition to some uncertainty about whether a post would be positive or negative, we also eliminated the most extreme reviews so as to avoid making the reviews' valence highly salient. Nonetheless, all posts were, in fact, either clearly negative or clearly positive.

As in the two earlier studies, five posts were available at each of the four simulated websites. Because there was a post-choice phase, we wanted the pre-choice phase to be similar in length for all subjects. To this end, we limited subjects to exactly five searches before the buy/not-buy decision. We also limited them to exactly three searches after their decisions.

Figure 3: Pre-choice and post-choice percentages of positive and negative information selected by subjects who were leaning against or toward buying the product (Study 3).



Subjects Sixty subjects were recruited from the pool of respondents provided by Amazon's Mechanical Turk. All completed the study online and were compensated for their time with \$1. We note that there is considerable satisfaction with this subject pool's convenience and low cost and rather few concerns about its representativeness or validity (e.g., Buhrmeister, Kwang, & Gosling, 2011; Paolacci, Chandler, & Ipeirotis, 2010; Sproue, 2011).⁵

Procedure. The procedure was identical to that of the two earlier studies, except for the unbiased values of each review. Instead of assuming equally spaced values, we ran a separate control group of 50 subjects, who independently rated the valence of every user review on the -50 to 50 scale. The means of these control ratings were used as the estimates of the true unbiased values of the various reviews and were entered into the computations of information distortion.

4.2 Results

Biased search We calculated the proportions of information selections that were from positive sources, both pre-choice and post-choice. These proportions are represented in Figure 3, split by prior leaning. Overall, and in contrast to Study 1 and 2, subjects exhibited a preference for negative over positive information pre-choice. This is explained by the popularity of the information coming from the online retailer (chosen 44% of the time pre-choice). Despite the fact that it was labeled as con-

⁵Buhrmeister et al. note that "the data obtained are at least as reliable as those obtained via traditional methods (p. 3)". One reported inferiority of the Mechanical Turk subject pool comes from Sproue's comparison with the traditional experimenter-controlled laboratory environment: "The results suggest that aside from a slightly higher subject rejection rates, AMT data are almost indistinguishable from laboratory data" (p. 155). Paolacci et al. also report higher incompleteness rates for Mechanical Turk subjects compared to student pools and laboratory settings, but still consider both rates acceptably high (both > 90%).

Table 3a: Effect of prior leaning on the selection of product information in Study 3, for log odds of Positive/Negative (unstandardized regression coefficients).

| Variables | Complete model | Pre-choice only | Post-choice only |
|--------------------------|----------------|-----------------|------------------|
| Intercept | -.45*** | -.54*** | -.37* |
| Prior leaning | 0.007 | -0.002 | .017* |
| Pre-Post(a) | 0.09 | | |
| Prior leaning x Pre-Post | .009* | | |

* $p < .05$, ** $p < .01$, *** $p < .001$; (a) Pre = -1 and Post = 1.

taining negative information about the product, subjects may have placed more trust in information from a retailer website than from a blog, a brand page, or a fan page. However, post-choice this preference for negative information reversed for subjects who decided to buy the product (see Figure 3). The pattern of information selected post-choice conformed to the selective exposure bias uniformly observed.

Once again we tested for an effect of prior leaning on information selection. Note that all after-choice leanings are based on the final choice confidence, signed positively if the subject chose to buy, and signed negatively if the subject chose not to buy. Therefore, the leaning coefficient we report was recorded on a scale from -50 (strong confidence toward not buying) to 50 (strong confidence toward buying), both pre-choice and post-choice. To perform this analysis, we ran a mixed logit model in SPSS, with the log odds of choosing positive information over negative information as the dependent variable and prior leaning, phase of the choice (Pre-Post), and their interaction as predictors. Subject ID and its interaction with Pre-Post were specified as random effects to control for within-subject effects. To be thorough, we also ran two separate logit models, one for information selection pre-choice and the other for information selection post-choice. Those two additional models were used to assess the significance of the slope of prior leaning. These three models are displayed in Table 3a.

The complete model showed that the interaction between prior leaning and Pre-Post was significant ($b = .009, p = .03$), indicating that the impact of leaning on information selection differed pre-choice and post-choice. The results of the pre-choice regression confirmed that leaning did not impact information search ($b = -.002, p = .59$) before a decision has been reached. Post-choice however, leaning significantly impacted the selection of information ($b = .017, p = .011$). Once a decision had been reached, a 10-unit increase in leaning increased the odds of selecting positive information over negative information by a factor of $\exp(10 \times .017) = 1.185$, an 18.5% increase.

To more precisely examine the impact of prior prefer-

ences on the valence of the information selected, we also computed the odds ratio of selecting positive information over negative information based on prior leaning towards buying (buy or not buy), both pre-choice and post-choice. Pre-choice, this odds ratio was 1.08, based on 38% of positive information selected when subjects lean towards buying vs. 36% when they lean towards not buying ($1.08 = [0.38/(1-0.38)]/[0.36/(1-0.36)]$). Post-choice, this odds ratio rose to 3.09 (57% of positive information for subjects who decided to buy the product vs. 30% for subjects who decided not to buy). Thus the odds of choosing confirming over disconfirming information clearly supported the presence of substantial selective exposure post-choice and little or none pre-choice.

Biased evaluation The Pre-Post contrast provided the opportunity to compare the before and after magnitudes of the preference-confirming bias in the evaluation of acquired information. Post-choice distortion was found to be lower (3.91 units) than pre-choice distortion (6.09), as in prior studies on distortion (Russo et al., 1998). However, this difference did not reach significance ($t(59) = -.99, p = .33$). Again, we ran a regression analysis in which information evaluation was the response variable, controlling for the estimated unbiased information evaluation means of each review that were obtained from the control group. In the complete model, the effect of prior leaning was positive and significant ($b = .20, p < .0001$). The interaction between leaning and Pre-Post did not reach significance ($p = .21$), signifying that the impact of prior leaning on information evaluation did not significantly differ pre-choice and post-choice. Two separate regressions for information evaluation pre-choice and post-choice confirmed that the pre-choice and post-choice slopes of leaning on information selected were similar (.20 vs. .35, both $ps = .0003$). The results of those regressions are displayed in Table 3b.

Desirability Just as in both earlier studies, we tested for subjects' desirability of searching for confirming information. Subjects' ratings of the desirability of the three "decision strategies" exhibited a similar pattern. The two

Table 3b: Effect of prior leaning on the evaluation of product information in Study 3 (unstandardized regression coefficients).

| Variables | Complete model | Pre-choice only | Post-choice only |
|--------------------------------|----------------|-----------------|------------------|
| Intercept | 0.77 | 0.11 | 1.37 |
| Prior leaning | .20*** | .20** | .35** |
| Pre-Post(a) | 0.7 | | |
| Prior leaning x Pre-Post | 0.05 | | |
| Mean valence of the review (b) | .85*** | .84*** | .82*** |

* $p < .05$, ** $p < .01$, *** $p < .001$; (a) Pre = -1 and Post = 1; (b) As assessed in our pre-test.

confirming biases were both rated as undesirable on average (for search, 2.72; for evaluation, 2.70), and were significantly lower than the scale midpoint (for search $t(59) = -2.14, p = .037$; for evaluation, $t(59) = -2.56, p = .013$). In contrast, the mean desirability rating of disconfirming search was 3.27, significantly higher both than the midpoint ($t(59) = 2.02, p = .048$) and than the desirability rating of confirming search ($t(59) = -2.84, p = .006$).

Search strategies We next analyzed subjects' reports of their own search strategies, using the same procedure and same three coders (initial agreement = .77) as in the two earlier studies. Of the 60 subjects, 51 gave codable strategies, and 16 of those fell in the catch-all category. Of the 35 remaining, 29 were classified as balanced and the last 6 as seeking more negative reviews. No subject described a search strategy that favored positive reviews.

4.3 Discussion

The above results replicated those of Studies 1 and 2 in (a) the absence of a pre-choice bias in search to support the currently preferred alternative, (b) the rated undesirability of such a bias, and (c) the overwhelming reporting of a balanced search strategy. Experiment 3 also verified the presence of the expected post-choice search bias to support the chosen alternative. One conclusion from this last result is that the absence of a pre-choice search bias in all three studies was not likely to be due to our using a method that was so artificial or insensitive that a search bias would not be observed under any circumstances.

5 General discussion

Three studies consistently found a complete absence of a pre-choice bias toward searching for preference-supporting information. The absence of this confirming search bias (known originally as the selective exposure to

information) occurred for products that were both hedonic and utilitarian, both expensive and inexpensive, and both high and low in expected brand loyalty. In addition to the absence of an effect of prior preferences on information selection, subjects' self-reported search strategies exhibited a clear tendency toward a balance of positive and negative information. Over all three studies, the total frequencies of balanced, positive, and negative strategies were, in order, 127, 1, and 20.

We have also attempted to place the familiar bias of selective exposure to information in the larger context of information use. For the present work that has meant testing for selective exposure pre-choice and verifying its expected presence post-choice. It has also meant assessing the second process associated with information, its evaluation. Again we confirmed the presence and assessed the magnitude of biased information evaluation (information distortion) both pre- and post-choice.

5.1 The cause of biased search

There is long-standing agreement that the main driver of the post-choice search bias to support the chosen option is the desire to defend that decision (e.g., Hart et al., 2009). In contrast, the presence or absence of the pre-choice equivalent of this search bias reflects the conflict between the goal of defending the current (tentative) preference and the goal of accuracy achieved by unbiased search. Our results are compatible with the dominance of the accuracy goal over the defense goal in the pre-choice phase of decision making. Our subjects partly supported that explanation with an average rating of the search for confirming information as undesirable. Thus, the overall pattern of results suggests that the absence of a preference-supporting search bias pre-choice is driven by the risk to a good decision from biasing information search.

Why, then, do most of the other studies of pre-choice selective exposure report its presence (viz., Fischer et al., 2011; Fraser-Mackenzie & Dror, 2009; Young et al., 2011)? The relative force of the accuracy goal should

have had just as much influence in these studies. Without more empirical work, we can only speculate that the method for extracting the currently preferred option may have contributed to the observed difference in findings. Specifically, in the three studies cited above, decision makers were instructed to state their tentative preference or in the work by Young et al., to have an existing tentative position. In contrast, Carlson and Guha (2011), who found no pre-choice selective exposure, used a horse-race metaphor that may have reduced the level of commitment to the tentative preference. In a similar way, the three studies reported above used an even less obtrusive method for obtaining the current preference and found no indication of this confirming bias. Until additional studies resolve the effect of the method for eliciting the current preference, we can only acknowledge its possible role in the observed differences in results across studies.

5.2 Limitations

Three limitations of our work should be noted. The first is the absence of a causal test of the impact of the accuracy goal on information search and evaluation. To overcome this limitation, future research could prime the goal of defense and see whether this priming increases the impact of leaning on information selection in the pre-choice phase. Alternatively, recording verbal protocols during the choice process may help researchers better understand how subjects reason during the task. Such protocols have already been used successfully with information distortion (Carlson, 2001; Carlson, Tanner, Meloy, & Russo, 2012) and could therefore also be applied to selective exposure.

Second, because we do not directly prime the goals of accuracy and defense, there remain other plausible explanations for our findings. In particular, other research has identified additional factors that may greatly influence the preference for decision-consistent information (e.g., Fischer, Schulz-Hardt, & Frey, 2008). For instance, when the amount of information from which people can select is limited, selective exposure has been shown to increase (Fischer, Jonas, Frey, & Schulz-Hardt, 2005). Thus, it is possible that we do not observe selective exposure in Studies 1 and 2 because the subjects' search for user reviews may have covered a large number of product attributes. Alternatively, the limited labeling of our reviews (number of stars) may not have communicated expected information quality; and expected quality has been shown to be a factor that influences selective exposure (Fischer, Jonas, Frey, & Schulz-Hardt, 2005; Fischer, Schulz-Hardt, & Frey, 2008). Although our Study 3 tried to assess selective exposure in conditions that were similar pre- and post-choice, we cannot dismiss the alternative explanations of amount and quality of information

in Studies 1 and 2.

The final limitation of our work is its generalizability. One of the dangers of experimental work is an overly general conclusion, even when the results are similar over different studies and consistent with a plausible explanation. We note first that all of our decisions were hypothetical. In actual decisions, might we observe at least some of the potential differences that we noted between hedonic and utilitarian products, between expensive and inexpensive ones, or between those enjoying high versus low brand loyalty? For instance, might pre-choice search be biased when actually purchasing a hedonic product with which a buyer would closely identify? Our results do not exclude the possibility that in some real decisions the goal of defending an emerging preference may be strong enough to compete with the accuracy goal and lead to selective exposure.

If the pre-choice bias toward confirming search is observed in real decisions, or at least in sufficiently realistic ones, how strong is its effect relative to other choice desiderata? That is, besides the question of its existence, is the pre-choice search for information influenced enough by a confirmation bias to overcome other desired criteria like the quality of the information, the credibility of its source, or its diagnostic value? At the same time that valid laboratory work tends to hold such other factors constant, they may play a substantial role in real decisions and influence the generalizability of any laboratory findings.

A final point regarding generalizability is the possibility of individual differences that affect susceptibility to selective exposure. Hart et al.'s (2009) meta-analysis found closed-mindedness to be a significant moderator of the extent of selective exposure. Yoon, Sarial-Abi, and Gurhan-Canli (2012) report that regulatory focus also influences selective exposure. These individual differences suggest one more reason why it is unlikely that selective exposure either always or never occurs.

5.3 Conclusion

Respecting the limits of the present studies and the risks of overly general claims, we nonetheless suggest that a pre-choice bias toward confirming search is unlikely to occur as long as two conditions are met. First, decision makers can tell when new information is likely to be confirming or disconfirming. Second, they recognize the risk to accuracy entailed by biasing their search toward confirming their currently leading alternative. Said differently, we suggest the accuracy goal will predominate whenever it can. However, and respecting the above caveats about overgeneralization, we fully acknowledge that any such claims must be empirically verified and subsequently qualified.

References

- Betsch, T., Haberstroh, S., Glöckner, A., Haar, T., & Fiedler, K. (2001). The effects of routine strength on adaptation and information search in recurrent decision-making. *Organizational Behavior and Human Decision Processes*, *84*, 23–53.
- Brownstein, A. L. (2003). Biased predecision processing. *Psychological Bulletin*, *129*, 545–568.
- Buhrmeister, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, *6*, 3–5.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–459.
- Carlson, K. A. (2001). In search of consumer goals. Unpublished doctoral dissertation, Cornell University.
- Carlson, K. A., & Guha, A. (2011). Leader-focused search: The impact of an emerging preference on information search. *Organizational Behavior and Human Decision Processes*, *115*, 133–141.
- Carlson, K. A., & Pearo, L. K. (2004). Limiting Predecisional Distortion by Prior Valuation of Attribute Components. *Organizational Behavior and Human Decision Processes*, *94*, 48–59.
- Carlson, K. A., Tanner, R. J., Meloy, M. G., & Russo, J. E. (2012). *Catching goals in the act of decision making*. Working Paper, Georgetown University, November.
- DeKay, M. L., Patiño-Echeverri, D., & Fischbeck, P. S. (2009). Distortion of probability and outcome information in risky decisions. *Organizational Behavior and Human Decision Processes*, *109*, 79–92.
- Engel, C., & Glöckner, A. (2013). Role-induced bias in court: An experimental analysis. *Journal of Behavioral Decision Making*, *24*, 272–284.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.
- Fischer, P., Jonas, E., Frey D., & Schulz-Hardt, S. (2005). Selective exposure to information: The impact of information limits. *European Journal of Social Psychology*, *35*, 469–492.
- Fischer, P., Schulz-Hardt, S., & Frey, D. (2008). Selective exposure and information quantity: How different information quantities moderate decision makers' preferences for consistent and inconsistent information. *Journal of Personality and Social Psychology*, *94*, 231–244.
- Fischer, P., & Greitemeyer, T. (2010). A new look at selective-exposure effects. *Psychological Science*, *19*, 384–389.
- Fischer, P., Lea, S., Kastenmuller, A., Greitemeyer, T., Fischer, J., & Frey, D. (2011). The process of selective exposure: Why confirmatory information search weakens over time. *Organizational Behavior and Human Decision Processes*, *114*, 37–48.
- Fraser-Mackenzie, P., & Dror, I. E. (2009). Selective information sampling: Cognitive coherence in evaluation of a novel item. *Judgment and Decision Making*, *4*, 307–316.
- Glöckner, A., & Betsch, T. (2008). Modeling option and strategy choices with connectionist networks: Towards an integrative model of automatic and deliberative decision making. *Judgment and Decision Making*, *3*, 215–228.
- Glöckner, A., & Herbold, A. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *Journal of Behavioral Decision Making*, *24*, 71–98.
- Hart, W., Albarracin, D., Eagly, A., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin*, *135*, 555–588.
- Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, *128*, 3–31.
- Hyman, H. H., & Sheatsley, P. B. (1947). Some reasons why information campaigns fail. *Public Opinion Quarterly*, *11*, 412–423.
- Keren, G., & Schul, Y. (2009). Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on Psychological Science*, *4*, 533–550.
- Lazarsfeld, P., Berelson, B., & Gaudet, H. (1944). *The people's choice: How the voter makes up his mind in a presidential campaign*. New York: Columbia University Press.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, *37*, 2098–2109.
- Meloy, M. G., & Russo, J. E. (2004). Binary choice under instructions to select versus reject. *Organizational Behavior and Human Decision Processes*, *93*, 114–128.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, *5*, 411–419.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
- Read, S. J., & Simon, D. (2012). Parallel constraint satisfaction as a mechanism for cognitive consistency. In B. Gawronski & F. Strack (Eds.), *Cognitive consistency*:

- A fundamental principle in social cognition.* (pp. 66-86). New York: Guilford Press.
- Russo, J. E. (in press). The predecisional distortion of information. In V. Reyna and E. Wilhelms (Eds.), *Neuroeconomics, judgment, and decision making*. New York: Psychology Press.
- Russo, J. E., Carlson, K. A., & Meloy, M. G. (2006). Choosing an inferior alternative. *Psychological Science, 17*, 899-904.
- Russo, J. E., Carlson K. A., Meloy M. G., & Yong K. (2008). The goal of consistency as a cause of information distortion (2008). *Journal of Experimental Psychology: General, 137*, 456-470.
- Russo, J. E., Meloy, M. G., & Medvec, V. H. (1998). Predecisional distortion of product information. *Journal of Marketing Research, 35*, 438-452.
- Russo, J. E., Medvec, V. H., & Meloy M. G. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes, 66*, 102-111.
- Simon, D., Pham, L. B., Le, Q. A., & Holyoak, K. J. (2001). The emergence of coherence over the course of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 1250-1260.
- Svenson, O. (1992). Differentiation and consolidation theory of human decision making: A frame of reference for the study of pre- and post-decision processes. *Acta Psychologica, 80*, 143-168.
- Svenson, O. (2006). Pre- and post-decision construction of preferences: Differentiation and consolidation. In S. Lichtenstein and P. Slovic (Eds.), *The construction of preferences* (pp. 356-371). Cambridge: Cambridge University Press.
- Svenson, O., & Benthorn, L. J. (1992). Consolidation processes in decision making: Post-decision changes in attractiveness of alternatives. *Journal of Economic Psychology, 13*, 315-327.
- Svenson, O., Salo, I., & Lindholm, T. (2009). Post-decision consolidation and distortion of facts. *Judgment and Decision Making, 4*, 397-407.
- Sprouse, J. (2011). A validation of Amazon Mechanical Turk for the collection of acceptability judgments in linguistic theory. *Behavior Research Methods, 43*, 155-167.
- Wilks, T. J. (2002). Predecisional distortion of evidence as a consequence of real-time audit review. *The Accounting Review, 77*, 51-71.
- Yoon, Y., Sarial-Abi, G., & Gurhan-Canli, Z. (2012). Effect of regulatory focus on selective information processing. *Journal of Consumer Research, 39*, 93-110.
- Young, M. J., Tiedens, L. Z., Jung, H., & Tsai, M. (2011). Mad enough to see the other side: Anger and the search for disconfirming information. *Cognition and Emotion, 25*, 10-21.