

A Satisficing Choice Model

June 28, 2011

Peter Stüttgen, Peter Boatwright, and Robert T. Monroe

Abstract

While the assumption of utility-maximizing consumers has been challenged for decades, empirical applications of alternative choice rules are still very recent. We add to these growing body of literature by proposing a model based on Simon’s idea of a “satisficing” decision maker. In contrast to previous models (including recent models implementing alternative choice rules), satisficing depends on the order in which alternatives are evaluated. We therefore conduct a visual conjoint experiment to collect search and choice data. We model search and choice jointly and allow for interdependence between them. The choice rule incorporates a conjunctive rule and, contrary to most previous models, does not rely on compensatory tradeoffs at all. The results strongly support the proposed model. We find that search is indeed influenced by product evaluations. More importantly, the model results strongly support the satisficing stopping rule. Finally, we discuss the different nature of choice predictions for the satisficing model and for a standard choice model and show how the satisficing model results in predictions that are more useful to retailers.

Keywords: Non-Compensatory Choice, Eye-Tracking, Visual Conjoint Experiment

1 Introduction

The large majority of choice models in the marketing literature focus on understanding the influence of product attributes or marketing variables on consumer choice. Implicitly, these models assume that the consumer has all information, or at least enough information to form a consideration set according to some rule. Yet, in reality a consumer needs to *acquire* information first before that information can, in turn, influence his decision. Thus, Pieters and Warlop (1999) suggest that visual attention may help understand consumer choice. Using eye-tracking methodology, they find that consumers do not gather all information about the alternatives. Moreover, they show that the percentage of people that look at the different alternatives is predictive of the alternatives' choice shares. In addition, they find directional support that skipping pieces of information (brand information, ingredient information, and/or the pictorial) for an alternative is negatively related to its choice share. In our data, the percentage of people who skipped some information for an alternative is an even better predictor of choice shares (Pearson's $r = -.75$) than the percentage of people who looked at a given alternative ($r = .36$). As discussed in more detail in our review of the literature, the phenomenon of skipped information within a product cannot be satisfactorily explained by standard utility maximizing models, even if allowing for search cost. Yet, the high correlation between choice shares and percentage of skipped information suggests that being able to understand *why* people skip some attribute information may be important for understanding consumer choice behavior.

The assumption that consumers are utility-maximizers has been criticized for a long time on grounds of the unrealistically high cognitive burden these rules impose on the decision maker (Simon 1955) as well as due to common violations of basic axioms of utility maximizing rules (e.g., transitivity; Loomes et al. 1991). Instead, consumers are believed to use simplified choice heuristics when making their choices (Gigerenzer and Todd 1999). While several simple alternative choice rules were proposed several decades ago (e.g., Coombs 1951), such rules have only been incorporated into empirical choice model very recently (e.g., Gilbride and Allenby 2004). We contribute to this new stream of research in the marketing literature by proposing a choice model based on Herb Simon's idea of "satisficing" choice (Simon 1955).

Satisficing is a simple choice rule in which the first alternative that is “good enough” according to some criterion is chosen. The outcome of a satisficing choice then is search path dependent. Moreover, the search path may be dependent on what “good enough” means. This interdependence between search and choice allows us to parsimoniously explain the phenomenon of skipped information in the framework of a simple choice rule. To do so, we model the search path and the resulting choice jointly. To gather the necessary data, we conduct a visual conjoint experiment. Using realistic stimuli in a conjoint design, we collect search path information using eye-tracking technology as well as the resulting choice data.

The satisficing choice rule implies a distinct stopping rule for the search process. Our results strongly confirm this stopping rule, providing evidence that consumers may in fact use a satisficing choice rule rather than a utility maximizing model in the product category that we studied. The satisficing choice rule is further supported by the results of a holdout prediction task in which the proposed satisficing model comfortably outpredicts a multinomial logit model.

The remainder of the paper is organized as follows: We will first briefly review the relevant streams of literature, then describe the experiment and the data, before explaining the proposed model and estimation. Finally, we present the results and conclude with a general discussion.

2 Literature Review

We first review the traditional approach to choice models and its merits but also argue why that approach cannot convincingly explain the observed data with skipped information. Then we review the literature on bounded rationality, and show that a satisficing choice rule in contrast can explain the observed data very parsimoniously. Finally, given our use of an eye-tracking experiment, we briefly review the relevant work on eye-tracking based search models.

2.1 Maximizing Choice

At least since Guadagni and Little (1983) introduced the multinomial logit model to the marketing literature in their seminal paper, the idea of a compensatory utility maximizing choice has been the predominant framework for empirical analyses of consumer choice. The theory of utility maximizing choice has its foundations in the tenets of microeconomics. Typically, utility is specified as a linear combination of the alternative's attributes, thereby making it a compensatory process (i.e., a "bad" value for one attribute can be compensated for by a "good" value for another attribute). The approach has proven to be straightforward to implement yet to yield valuable managerial insights, for instance enabling managers to segment the market (e.g., Kamakura and Russell 1989) or to understand the impact of marketing decisions (e.g., Gupta 1988). Building on this framework, the more recent advent of structural models has allowed researchers to examine consumers' strategic and forward looking behavior (e.g., Sun 2005).

However, the assumption of a rational consumer with unlimited cognitive capabilities, as theoretically appealing as it may be from a normative standpoint, has long been challenged as an appropriate representation of actual human decision makers (e.g., Simon 1955; Kahneman and Tversky 1979). Even proponents of the utility maximizing approach typically agree that decision makers may not actually make decisions following the rules of the model, but rather they act *as if* they did. The models then are seen as a description of the outcome rather than the process. Nonetheless, following Shugan's call for incorporating a "cost of thinking" (Shugan 1980) to allow for more realistic models, the literature on choice models has started to account for limited consumer search and introduced "cognitively less demanding" (for the decision maker) models .

Importantly, Hauser and Wernerfelt (1990) and Roberts and Lattin (1991) incorporated consideration sets into choice models, i.e. they proposed a two stage process in which only a subset of the available alternatives is selected in the first stage for a utility maximizing choice in the second stage. The formation of the consideration set was originally dependent on cost-benefit tradeoffs for including an additional brand into the consideration set, or, with the rise of structural models, an explicit tradeoff of search cost and expected benefit

(Mehta et al. 2003). Yet, these models of “constrained utility maximization” typically in fact *increase* the computational burden of the decision maker rather than decrease it. In Gigerenzer and Todd (1999)’s words, “[t]he paradoxical approach of optimization under constraints [i.e., optimization including a search or other sort of cost] is to model “limited” search by assuming that the mind has essentially unlimited time and knowledge with which to evaluate the costs and benefits to further information search” (p. 11).

Yet, even if we embrace constrained utility maximization as an appropriate framework to model consumer decision making, we encounter serious problems in trying to explain the search patterns observed in our visual conjoint experiment (to be described in section 3). Incorporating cost of thinking (i.e., search cost and or cost for evaluating information) in the model framework can easily explain why consumers may not evaluate all available alternatives. If consumers use a stopping rule for search based on cost-benefit tradeoffs, it may be optimal to not search all available options. However, search costs do not plausibly explain why consumers may start evaluating some alternatives but not collect all information about them, a frequent pattern in the search process. One explanation might be that search cost within a product is higher than across products. However, this does not seem likely since (a) search cost should be very low within a product as the shopper only needs to move his eyes minimally, and (b) integration of information should not be very difficult since the product category (instant [Ramen] noodles) does not involve difficult tradeoffs between attributes. Alternatively, incomplete search within a product could occur in a search cost framework if a given consumer cares a lot more about one attribute (say, flavor) than another (say, price). In that case, knowing that a product has a flavor he really dislikes may make it unprofitable for him to acquire the other attributes for this product, as even the lowest price may not offset the disutility caused by the flavor. However, in our data we find that it is not always the same attribute that is missing (within a person), so therefore this cannot be the correct explanation.

As we will show below, though, both limited search across alternatives as well as limited search within alternatives can be easily explained by a model that is not based on a utility maximizing framework but instead uses a satisficing choice rule.

2.2 Bounded Rationality

In his above mentioned critique of the rational utility-maximizing agent, Simon says “the task is to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information to the computational capacities that are actually possessed by [...] man” (Simon 1955, p. 99).¹ This is the foundation of what has become to be known as “bounded rationality”. In this view, decision makers are aware of their cognitive limits and therefore rely on simplified choice rules (depending on the task). The best-known of these simplified choice rules are the lexicographic rule, the conjunctive and disjunctive rules, and elimination by aspects.

In the lexicographic rule (von Neumann and Morgenstern 1947), a decision maker focuses on the attribute that is most important to her and simply chooses the alternative that is best on that particular attribute. If there is a tie, she compares the tied alternatives on her second most important attribute and chooses the alternative that is preferred according to that attribute. The process continues until a unique choice is found (or until all attributes are exhausted).

In the conjunctive and disjunctive rules (Coombs 1951; Dawes 1964), the decision maker has individual threshold levels for all attributes. In the conjunctive rule, every product that passes *all* of these thresholds is acceptable to him, whereas in the disjunctive rule all products that pass *at least one* threshold are acceptable.

Finally, elimination by aspects (Tversky 1972) essentially is a combination of the lexicographic and the conjunctive rules in which a decision maker first focuses on the most important attribute, but not only the best alternative makes it to the second round, but all alternatives passing the threshold for the particular attribute.

Notice that all of these choice rules are non-compensatory, i.e. a “bad” value for one attribute may be enough for not choosing a particular product, irrespective on how good it may be on other attributes. Thus, these decision rules tremendously simplify the decision process, as the decision maker does not have to evaluate any tradeoffs between attributes.

¹Following the rules of the well-known children’s game, we attempt to do exactly as ‘Simon says’ in this paper.

Despite the fact that these models were proposed several decades ago, and despite more and more behavioral evidence that central assumptions of the utility maximizing framework seem to be violated in reality (see Bettman et al. 1991 for a review of consumer decision making), most empirical applications of non-compensatory models stem only from the last decade.

Several of these applications have extended the linear utility framework to be able to capture screening rules based on these simplified rules (e.g., Swait 2001; Elrod et al. 2004), while others have directly modeled the simplified rules (see Gilbride and Allenby 2004 and Jedidi and Kohli 2005 for the conjunctive and disjunctive rules, Kohli and Jedidi 2007 for the lexicographic rule, and Gilbride and Allenby 2006 for elimination by aspects). However, note that these simplified choice rules (1) assume that the decision maker knows the values of at least one attribute for *all* alternatives and (2) that the conjunctive and disjunctive rule only result in a set of acceptable set of products. Thus, the disjunctive and conjunctive rules so far have either been used to predict a whole set (acceptable MBA candidates, Jedidi and Kohli 2005) or are followed by or combined with a compensatory choice rule (Swait 2001; Elrod et al. 2004; Gilbride and Allenby 2004).

“Satisficing”, a term coined by Simon combining “satisfactory” and “sufficing”, in contrast results in a unique choice *and* can explain limited search. The decision process is very simple: Start by evaluating one alternative. If it is satisfactory (according to a criterion to be defined), choose that product and stop searching. If not, evaluate the next alternative. Continue this process until you have found a satisfactory alternative.

It is obvious why people following a satisficing choice rule may not search all alternatives. However, depending on the satisfaction criterion, satisficing can also explain incomplete information acquisition within a product. For instance, say the satisfaction criterion is given by a conjunctive rule. Then, once the decision maker knows that the product fails to meet the threshold on one attribute (be it price, flavor, or brand) there is no reason to continue the search within this product.

One difficulty for the empirical application of a satisficing rule is that the choice outcome depends on the sequence of evaluation. If there is more than one satisfactory product, the decision maker will choose whichever she comes across first. Thus, it is essential to know

the search sequence. We therefore use eye-tracking technology to record the sequence of information acquisition in our visual conjoint experiment.

2.3 Eye-Tracking

Eye-tracking hardware has improved tremendously over the last 15 years, allowing for unobtrusive observation of a person’s eye fixations. In an early application, Russo and Leclerc (1994) relied on human coders to code the location of the fixations based on a video of the participants’ face. They identified three processing stages in consumer choice: Orientation, evaluation, and verification. In the verification stage, most relevant to our work for reasons that will become clear later, participants continue to search and acquire information despite already having made a choice. In the study by Pieters and Warlop (1999) the data collection was automated, but consumers had to keep their heads fixed to the apparatus. Improvements in soft- and hardware nowadays allow participants to move freely in about a 25x25x25 inch box, while the location of their eye fixations is determined based on the eyes’ reflection of infrared signals.

Applications of eye-tracking in marketing research have brought valuable insights in consumers’ processing of print ads (e.g., Pieters et al. (2002)), the optimal design of TV commercials (Teixeira et al. 2010), and the effect of in-store marketing activities (Chandon et al. 2009), among others. For a review of the findings from eye-tracking applications in marketing, see Wedel and Pieters (2008). Most relevant to our application, though, is the research explaining consumer search patterns. van der Lans et al. (2008a) show that consumer search is influenced both by features of the stimulus, so-called bottom-up effects, and by strategic or intentional strategies, so-called top-down effects. Liechty et al. (2003) propose a hidden Markov Model to capture two distinct types of search, local search and global search. In the local search state “stimuli are explored in detail by extracting information from specific and adjacent locations”, whereas the global state “is characterized by longer saccades” (i.e., movements between fixations) and “stimuli are explored to identify locations to extract information” (Liechty et al. 2003, p. 520).

3 Visual Conjoint Experiment

As discussed above, the empirical application of a satisficing model requires knowledge of both search path and product choice. Moreover, since search and product choice may be inter-dependent, we model the search path and the choice jointly. To collect the data needed to do this, we conduct what we term a “visual conjoint experiment”, i.e., we develop a standard conjoint design, but then translate the resulting choice sets into realistic images of shelves from which participants make their selections.

3.1 Participants

The experiment was conducted at the Doha, Qatar campus of Carnegie Mellon University. Participants were 75 undergraduate students from the Doha, Qatar campuses of Carnegie Mellon University, Texas A&M University, Georgetown University, Northwestern University, and Cornell Medical College. 11 of these participants are excluded from the analysis due to calibration problems and/or incomplete eye-recordings, leaving a total of 64 students (29 female, 35 male) in the sample.² Participants’ age ranges from 17 to 23 years, with a mean of 19.88 years. Participants’ nationalities are predominantly (~55%) South Asian (e.g., Indian, Pakistani, Bangladeshi), and Middle Eastern countries combine for a total of 18 participants (28%).³ Six out of the 64 participants were U.S. American. Subjects were paid approximately \$14 (depending on their choices), and sessions lasted between 30 and 60 minutes.

3.2 Stimuli and Procedure

We choose instant noodles (also known as “Ramen noodles”) as the product category. Products vary on price, flavor, and brand. We use four brands, five equidistant price levels (ranging from ~\$1.10 to ~\$1.90 for a five pack of noodles), and ten flavors. The brands and flavors were selected from brands and flavors present in the local market. Similarly, the

²The calibration procedure is explained in the following subsection.

³Many of the participants have lived in Qatar for most, if not all, of their lives.

price levels span the price range found in the local market. The conjoint design consists of 15 choice sets with 15 alternatives each. We translate each choice set into an image of three shelves with five alternatives each. To approximate a realistic amount of clutter on the shelves, each alternative has four facings. See Figure 1 for an example. We used a 50 inch HD television (1920 x 1080 pixels) in the experiment, which allowed for the products to be approximately real-life sized and made all information easily readable.



Figure 1: Example stimulus

Subjects participated in the experiment in individual sessions. After reading the instructions, including a list of the available brands and flavors as well as an example stimulus, the eye-tracking software was calibrated. For the individual specific calibration, subjects were asked to follow a dot moving around the screen with their eyes to “teach” the software how eye movements relate to location on the screen. Calibration was repeated after one third and after two thirds of the experiment to ensure high quality data. After calibration, the first stimulus appeared on the screen and participants could take as long as they needed to make a decision. Once they reached a decision, they clicked a button on a presentation clicker which caused the screen to blur and the products were overlaid with letters from A

to O. This was done to prohibited acquisition of additional information after a choice has been made; note that Russo and Leclerc (1994)’s verification state, if present, then inherently becomes a part of the recorded search path. Subjects then indicated their choice by announcing the corresponding letter to the experimenter, or said that they chose not to buy anything from this particular choice set (Pieters and Warlop 1999). After the last choice, participants completed a questionnaire to collect, among other things, explicit measures of their preferences.

To ensure that the task was incentive compatible, one of the choice sets was selected at the end of the experiment and the corresponding purchase realized (i.e., participants received their chosen item and paid the respective price from their participation fee).

3.3 Data

For each participant, we then have the 15 choice outcomes, the questionnaire responses, as well as the sequence of the locations of eye fixations for each choice set. Since our interest lies mainly in information acquisition, we aggregate the pixel-level data into meaningful areas of interest (AOI), namely the price tag, the flavor information, and the rest of the package for each of the alternatives, plus fixations on the background (Pieters and Warlop 1999; Shi et al. 2010). Following Shi et al. (2010) we exclude fixations on the background as well as consecutive repeat fixations on the same AOI, as they are not informative about a consumers information acquisition process. Thus, we have 45 AOIs (15 products with 3 AOIs each) which provide an exhaustive and mutually exclusive partition of each stimulus. Since the packaging distinguishes brands and brands are well-known, we assume that participants learn a product’s brand by looking anywhere on the packaging (including the flavor AOI), whereas they have to fixate on the corresponding AOI to learn the flavor or price.

Figure 2 shows the average number of fixations per choice set. It is obvious that participants tend to search longer in the first few images, most likely to get used to the task. For the effect of number of fixations on the likelihood of termination (see section 4.2.3), we therefore normalize the number of fixations by the average number of fixations for the respective




Figure 2: Average Number of Fixations per Choice Set

Table 1: Brand Choice Shares

Fantastic	Indomie	Koka	Maggi
19.0%	22.4%	12.7%	45.9%

choice set.⁴ The number of fixations within a subject varies greatly across choice sets; even when only considering the last 11 choice sets (i.e., when average fixations have stabilized) the mean (across participants) standard deviation (for one participant across choice sets) is 13.9 fixations. This suggests that participants do not simply follow a fixed-search stopping rule, but employ a more variable stopping rule depending on the information acquired in a particular search.

Tables 1 and 2 provide a summary of brand and flavor choices, respectively, giving a first indication of consumer preferences. The Maggi brand as well as the chicken and onion chicken flavors are clear consumer favorites. Overall, participants decided not to buy in 7.0% of choices.

⁴Results are qualitatively equivalent if we instead exclude the first three choice sets. We prefer the normalization as to not lose the information contained for other parts of the model.

Table 2: Flavor Choice Shares

Beef	Cheese	Chicken	Curry	Lobster
6.6%	11.8%	23.1%	12.0%	4.7%
Mushroom	Onion Chicken	Shrimp	Tomato	Vegetable
4.6%	18.0%	6.9%	5.8%	6.5%

4 Model

For ease of presentation, we split the model into two interrelated parts: search and choice. Though we present them separately, satisficing truly combines both as the choice depends on the sequence and the search depends on previous judgments of satisfaction. To accommodate for the verification stage identified by Russo and Leclerc (1994), we relax Simon’s strict satisficing rule to allow for continued search after encountering the first satisfactory alternative. Nonetheless, even in the relaxed version the probability of stopping the search should increase significantly after encountering the first satisfactory alternative. The presence or absence of this effect can then be interpreted as an indicator for whether participants in fact used a satisficing choice rule or not.

4.1 Choice

We use a conjunctive rule as the satisfaction criterion to be used within the satisficing choice. That is, for each of the three attributes an individual has a set of acceptable levels, and only if she (1) has learned all three attribute levels for a given product and (2) all three attribute levels are acceptable, the product is judged to be satisfactory to her.⁵ In contrast, as soon as at least one unacceptable attribute level has been found for a product, the product is judged to be unsatisfactory. If no unacceptable attribute level has been encountered yet and not all attribute levels for a given product have been learned, the products status is undetermined. Note that this implies that all products are undetermined at the beginning of the search and

⁵To avoid confusion, we will use un/satisfactory for the product level and (not) acceptable for the attribute level for the remainder of the paper.

can change status at any point during the search.

More formally, let Γ_{iA} denote the acceptable set for attribute $A \in \{B,F,P\}$ (brand, flavor, and price) for individual i and I_{ijf} be the cumulative information set of individual i in choice set j as of (and including) fixation f . Further, let a represent the distinct levels of attribute A and let a_{jx} denote the particular attribute level of product x in choice set j . S_{ijxf} and U_{ijxf} are indicators for whether product x has been judged to be satisfactory or unsatisfactory, respectively, as of fixation f in choice set j by individual i .

Suppressing i and j for readability and using \mathcal{I} to represent the indicator function, we then have

$$S_{xf} = \prod_A \mathcal{I}_{a_x \in \Gamma_A} \cdot \mathcal{I}_{a_x \in I_f} \quad (1)$$

and

$$U_{xf} = \max_A (\mathcal{I}_{a_x \notin \Gamma_A} \cdot \mathcal{I}_{a_x \in I_f}) \quad (2)$$

The product over all attributes in equation 1 reflects the notion that *all* attributes of a product have to be known and acceptable for the product to be satisfactory. In contrast, a product is unsatisfactory if *at least one* attribute is known and unacceptable, as captured the maximum-operator in equation 2.

Since we allow for a verification stage in the search process, it is possible that a decision maker finds more than one satisfactory alternative before terminating the search. In order to keep the choice model as close to Simon’s original idea of satisficing, i.e., all a consumer cares about is passing a certain threshold rather than some relative ranking, we posit random choice between all satisfactory alternatives. We also introduce the notion of a “trembling hand” to allow for unobserved error. In game theory, a “trembling hand” allows for non-zero probabilities of actions off the equilibrium path (Osborne and Rubinstein 1994). In our context, it allows for non-zero choice probabilities for undetermined alternatives as well as for no-choice (which should in theory never be chosen in the presence of a satisfactory alternative). However, we posit a choice probability of exactly zero for all unsatisfactory options as well as all options which were never fixated on at all.

Let C_{ijx} be product x ’s choice index at the end of consumer i ’s search in choice set j .

We then let $C_{ijx} = 1$ for all products that are satisfactory at the end of the search, $C_{ijx} = 0$ for all products that are unsatisfactory at the end of the search and for all products never fixated on, and $C_{ijx} = \tau_i$, $0 < \tau_i < 1$ for all other products and for no-choice. Product x^* 's choice probability is then given by

$$Pr(x^*) = \frac{C_{x^*}}{\sum_x C_x}$$

Consumer heterogeneity in preferences is captured by a hierarchical structure. For brands and flavors, the individual level acceptability follows a Bernoulli distribution, i.e.,

$$\mathcal{I}_{a \in \Gamma_{iA}} \sim \text{Bern}(\hat{\gamma}_a) \text{ for } A \in \{B, F\} \quad (3)$$

For prices, we estimate the highest price acceptable (of the prices used in the experiment); the individual level threshold prices follow a multinomial distribution, i.e.,

$$\hat{p}_i \sim \text{MN}(1, \vec{p}) \quad (4)$$

and

$$p \in \Gamma_{iP} \text{ if and only if } p \leq \hat{p}_i \quad (5)$$

Finally, τ_i follows a Beta distribution, i.e.,

$$\tau_i \sim \text{Beta}(\hat{\tau}_1, \hat{\tau}_2)$$

4.2 Search

Following the literature based on Liechty et al. (2003), we model consumer information search with a modified hidden Markov model. There are two unobserved search states, namely local and global search, as well as a termination state, defined by the button press of the participant at the end of his search.

4.2.1 Global Search

Recall that global search consists mainly of large “jumps” across the stimulus to explore different areas (Liechty et al. 2003). Since this implies moving to an area that has previously been only in peripheral vision, targeting a specific location for the next fixation is difficult for the participant. Therefore, we assume product-level effects for the probabilities of moving, and the probability of the exact location within that product (i.e., which AOI) is proportional to the size of the respective AOI.

In the global state, eye movements are largely influenced by the saliency and/or luminance of image areas (van der Lans et al. 2008b). Since these are confounded with brand through brand-specific packaging, we use separate brand intercepts to capture this effect. Moreover, the status of the alternative may influence the probability of moving to a respective product. If a product has already been judged to be satisfactory or unsatisfactory, there is no reason from an information acquisition viewpoint to return to that product later in the search. Notice that by incorporating the status into the search probabilities, these probabilities become path-dependent and vary with time.

Let AOIs be denoted by h . Further, let $b_j(h)$ and $p_j(h)$ be the brand and product that h belongs to in choice set j , respectively, and let $r_j(h)$ be the ratio of h size in choice set j relative to the size of $p_j(h)$. With this notation (but suppressing i and j again), the above considerations can be expressed as follows.

Let \vec{S}_{f-1} and \vec{U}_{f-1} be the vectors containing the all product-level indicators $S_{x,f-1}$ and $U_{x,f-1}$. The probability of moving to AOI h at fixation f in the global state g conditional on the statuses of all alternatives (summarizing the previous search) is then given by

$$\eta_{fg}(h|\vec{S}_{f-1},\vec{U}_{f-1}) = \frac{\psi_{fg}(h|\cdot)}{\Psi_{fg}}$$

where

$$\psi_{fg}(h|\cdot) = r(h) \cdot \exp(\phi_{g0b(h)} + \phi_{g1}S_{p(h),f-1} + \phi_{g2}U_{p(h),f-1}) \quad (6)$$

and $\Psi_{fg} = \sum_h \psi_{fg}(h|\cdot)$ is the appropriate normalizing constant and ϕ s are parameters to be estimated.

4.2.2 Local Search

In contrast to global search, local search is aimed at gathering specific pieces of information in the same area of the image as the previous fixation (Liechty et al. 2003). Thus, re-fixating on the same product or moving to an adjacent product should be most likely. Moreover, staying in the same area of the image allows for targeted search for specific attribute information. In the local state, people may also use systematic search strategies such as search by attribute. Finally, similar to the global state, we expect people to be less likely to return to a product that has already been determined to be unsatisfactory.

Let $R_{ijf}(h|h_{ij,f-1})$, $N_{rijf}(h|h_{ij,f-1})$, and $N_{lijf}(h|h_{ij,f-1})$ be dummy variables for whether moving to AOI h would constitute a re-fixation on the same product, a move to the neighboring product to the right, or a move to the neighboring product to the left, respectively.⁶ In case the product last fixated on was on the left (right) edge of the shelf, $N_{lijf}(h)$ ($N_{rijf}(h)$) equals one for the AOIs corresponding to the product(s) directly above and/or below that product. Further, let $L_{ijf}(h|h_{ij,f-1})$ (or short $L_{ijf}(h)$) be the sum of these three indicators, thereby defining the “local” area around the last fixation. Finally, to keep track of the type of information represented by a given AOI, let $A(h)$ denote whether h is a price tag ($A(h) = P$), flavor information ($A(h) = F$), or the remaining packaging ($A(h) = B$). This helps to investigate target information search as well as search by attribute.

Similarly to the specification for the local state and once again suppressing i and j , we then have

$$\eta_{fl}(h|h_{f-1}, \vec{U}_{f-1}) = \frac{\psi_{fl}(h|\cdot)}{\Psi_{fl}}$$

where

$$\psi_{fl}(h|\cdot) = \exp \left(\begin{array}{l} \phi_{l0}L_f(h) + \phi_{l1}R_f(h) + \phi_{l2}N_{rf}(h) + \phi_{l3}U_{p(h),f-1} \\ + \sum_A[\phi_{l4A}\mathcal{I}_{A(h)=A} \cdot L_f(h)] + \sum_A[\phi_{l5A}\mathcal{I}_{A(h)=A}L_f(h) \cdot \mathcal{I}_{A(h)=A(h_{f-1})}] \end{array} \right) \quad (7)$$

and $\Psi_{fl} = \sum_h \psi_{fl}(h|\cdot)$ again is the appropriate normalizing constant. ϕ_{l0} to ϕ_{l2} represent

⁶For readability, we will omit the conditioning term for the remainder of the paper as it can be easily derived from the first set of subscripts.

the effects of staying in the same area as well as potential incremental effects of refixating on the same product or moving to the right (both relative to moving to the left). The second line captures the targeted search for information (ϕ_{l4A} only comes into play if moving to AOI h is a “local” move to attribute A) as well as the strategic search by attribute (ϕ_{l5A} only comes into play if moving to AOI h is a “local” move and both the current AOI and h correspond to attribute A).

4.2.3 Transition Probabilities

To complete the hidden Markov model for the search, we need to specify the transition probabilities between the states. Letting $s_{ijf} \in \{g, l, t\}$ be the state of the f th fixation of individual i in choice set j , the transition probabilities take the general form

$$Pr(s_{ijf} = s^* | f, s_{ij,f-1}, h_{ij,f-1}, \vec{S}_{ij,f-1}, \vec{U}_{ij,f-1}) = \frac{\pi_{ijf}(s^* | \cdot)}{\sum_s \pi_{ijf}(s | \cdot)} \quad (8)$$

For the transition to the global and to the local state, we again allow transitions to depend on the status of the product last fixated on. If the status is already determined to be either satisfactory or unsatisfactory, transition to the global state may be more likely in order to move to a different area of the image. Moreover, it should depend on the previous state, allowing for auto-correlation between states. Again suppressing i and j , we thus have

$$\pi_f(g | \cdot) = \exp(\lambda_{g0} + \lambda_{g1} \mathcal{I}_{s_{f-1}=g} + \lambda_{g2} S_{p(h_{f-1}),f-1} + \lambda_{g3} U_{p(h_{f-1}),f-1}) \quad (9)$$

and

$$\pi_f(l | \cdot) = \exp(\lambda_{l0} + \lambda_{l1} \mathcal{I}_{s_{f-1}=g} + \lambda_{l2} S_{p(h_{f-1}),f-1} + \lambda_{l3} U_{p(h_{f-1}),f-1}) \quad (10)$$

Most interestingly, though, explicitly modeling the transition to the termination state allows for better insights into what causes consumers to quit searching (Liechty et al. 2003). This is of particular interest in our application, since the satisficing choice rule has a very distinct stopping rule that we can directly model. Though the stopping rule is not deterministic due to allowing for the verification stage, satisficing implies that transitioning to

the termination state should be significantly more likely after finding the first satisfactory alternative. In addition, people probably are more likely to quit searching, the longer the search has been. With the understanding that the transition to the termination state occurs after the last recorded fixation, we then let

$$\pi_f(t|\cdot) = \exp(\lambda_{t0} + \lambda_{t1}\mathcal{I}_{s_{f-1}=g} + \lambda_{t2}f^* + \lambda_{t3}\max(\vec{S}_{f-1})) \quad (11)$$

where f_j^* is the running count of fixations normalized by the average number of fixations for the respective choice set, as explained in section 3.3. If consumers use a satisficing rule, λ_{t3} (capturing the implied stopping rule of the satisficing model) should be positive. Moreover, its magnitude relative to λ_{t2} determines how important that effect is relative to the effect of the number of fixations (which can be interpreted as a proxy for search cost and/or fatigue).

4.2.4 Heterogeneity

Similar to the choice part of the model, all parameters in the search part of the model allow for inter-individual heterogeneity through a normal hierarchical structure, i.e., we have

$$\phi_i. \sim N(\bar{\phi}., \sigma.)$$

and

$$\lambda_i. \sim N(\bar{\lambda}., \varsigma.)$$

5 Estimation and Identification

The model is completed by a set of uninformative priors (see appendix A.1 for details) and estimated with an MCMC algorithm in a Bayesian framework (Gelfand and Smith 1990; Casella and George 1992). The MCMC allows for efficient integration over the inherent discontinuities of the model caused by the indicator nature of acceptability and satisfaction judgments. We use 100,000 draws after discarding 30,000 draws as a burn-in period. Convergence is assessed using the the Heidelberger and Welch convergence diagnostic as

implemented in the **boa** package for R (Heidelberger and Welch 1983; Smith 2007).

Several restrictions have to be placed on the above model for identification. Analogous to standard choice models, the brand intercepts in equation 6 for the global search are not separately identified. We therefore normalize $\phi_{g0i, Fantastic}$ to zero. Equation 7 for the local search is also not uniquely identified. In particular, since the brand, flavor, and price AOIs collectively make up the complete product, ϕ_{l0i} is not separately identified from the set $\{\phi_{l4Bi}, \phi_{l4Fi}, \phi_{l4Pi}\}$ and only weakly identified from the set $\{\phi_{l5Bi}, \phi_{l5Fi}, \phi_{l5Pi}\}$ (recall that consecutive repeat fixations on the same AOI are excluded from the data and that $R.(h)$ therefore does not play into the systematic search by attribute). We therefore normalize ϕ_{l4B} and ϕ_{l5B} to zero. Moreover, the initial fixation in each search sequence is assumed to be in the global state and its location to be exogenous (van der Lans et al. 2008a).

Once again analogous to intercepts in standard models, only two out of the three λ_{s0} and the three λ_{s1} parameters for the transition probabilities are identified. (Note that $\lambda_{s0} + \lambda_{s1}$ is the intercept conditional on the last fixation being in the global state.) We thus normalize λ_{g0} and λ_{g1} to zero. The identification of λ_{g2} and λ_{l2} as well as of λ_{g3} and λ_{l3} relies on fitting the relative probability of termination. However, as reported below we find that the probability of termination is so low (at least before finding a satisfactory option) that only the difference of these parameter is well informed, as the termination probability changes only minimally for a wide range of absolute values of these parameters. We therefore set $\pi_{ijf}(g|\cdot) = 1$ for all i, j , and f . However, this still allows us to investigate the relative impact of an alternative’s status on transitioning to the global vs. the local state as well as, more importantly, the use of the stopping rule implied by the satisficing choice.

In addition to the “full model” described above, we also estimate an “independent model” in which all search model parameters corresponding to the status of the alternatives (ϕ_{g1i} , ϕ_{g2i} , ϕ_{l3i} , λ_{l2i} , λ_{l3i} , and λ_{l3i}) are set to zero. This allows us to analyze the influence of modeling the two parts jointly on the obtained results.

6 Results

Analog to the presentation of the model, we separate the presentation of the results into search and choice. We focus on the results of the full model and discuss the results of the independent model wherever they allow extra insight into the model.⁷ Finally, we present a holdout prediction analysis and compare the result against standard logit models.

While we will analyze all parameter results in detail, it is worth highlighting the results on the parameters corresponding to the alternatives' statuses. Since estimating these indicators and having both choice and search depend on them is one of the main features of the proposed model, it is important to check whether they do in fact affect search and choice in reasonable ways. As discussed below, we find that all of the parameters corresponding to the alternatives' statuses - including the one for the stopping rule - are non-zero (i.e., the 95% highest density interval of the posterior does not include zero) and all but one have the expected sign. The one exception is explained by a closer look at the data. Thus, the results suggest that these un/satisfactory judgments are in fact meaningful and have real impact on continued search.

6.1 Search

Table 3 presents the posterior means and standard deviations for the population level hierarchies for the search model parameters.

For the brand intercepts in the global search, it is interesting to compare them to the choice shares given in Table 1. Recall that packaging is confounded with the brands in our experiment, so one could expect the search parameters to reflect some of the brand preferences. However, excluding the most preferred brand of Maggi, the order of parameter estimates is actually reversed to the order of brand shares. This suggests that while colors (e.g., yellow for Maggi) may be useful for top-down search goals (van der Lans et al. 2008a), this process is far from perfect.

⁷None of the results of the independent model differ qualitatively from the full model unless mentioned.

Table 3: Posterior Means (and Standard Deviations) for the Search Parameters

Global Search					
$\bar{\phi}_{g0,Indomie}$	-.28	(.13)**	$\sigma_{g0,Indomie}$.80	(.20)
$\bar{\phi}_{g0,Koka}$.12	(.07)*	$\sigma_{g0,Koka}$.15	(.05)
$\bar{\phi}_{g0,Maggi}$.48	(.08)**	$\sigma_{g0,Maggi}$.32	(.08)
$\bar{\phi}_{g1}$	1.06	(.10)**	σ_{g1}	.44	(.13)
$\bar{\phi}_{g2}$	-.94	(.09)**	σ_{g2}	.25	(.09)
Local Search					
$\bar{\phi}_{l0}$	3.01	(.09)**	σ_{l0}	.43	(.09)
$\bar{\phi}_{l1}$	1.14	(.05)**	σ_{l1}	.12	(.03)
$\bar{\phi}_{l2}$.04	(.04)	σ_{l2}	.07	(.02)
$\bar{\phi}_{l3}$	-.62	(.04)**	σ_{l3}	.08	(.02)
$\bar{\phi}_{l4,Flavor}$	-.08	(.04)**	$\sigma_{l4,Flavor}$.09	(.02)
$\bar{\phi}_{l4,Price}$	-.77	(.06)**	$\sigma_{l4,Price}$.22	(.05)
$\bar{\phi}_{l5,Flavor}$	-.38	(.06)**	$\sigma_{l5,Flavor}$.12	(.03)
$\bar{\phi}_{l5,Price}$	1.14	(.08)**	$\sigma_{l5,Price}$.33	(.08)
Transition Probabilities					
λ_{l0}	1.54	(.12)**	ς_{l0}	.78	(.16)
λ_{l1}	-.08	(.14)	ς_{l1}	.88	(.23)
λ_{l2}	-.28	(.06)**	ς_{l2}	.08	(.02)
λ_{l3}	-.47	(.05)**	ς_{l3}	.07	(.02)
λ_{t0}	-5.17	(.21)**	ς_{t0}	1.26	(.35)
λ_{t1}	.11	(.10)	ς_{t1}	.18	(.08)
λ_{t2}	1.68	(.15)**	ς_{t2}	.64	(.24)
λ_{t3}	2.34	(.14)**	ς_{t3}	.41	(.17)

* 90% Highest Density Interval does not include zero

** 95% Highest Density Interval does not include zero

As expected, a consumer is less likely to return to an alternative already judged to be unsatisfactory ($\bar{\phi}_{g2} = -.94$). In contrast, a consumer is more likely to return to an alternative already judged to be satisfactory ($\bar{\phi}_{g1} = 1.06$). We did not expect this result, since there should be no reason to return to an already determined alternative from an information acquisition viewpoint. One possible explanation might be that consumers make explicit comparisons between different satisfactory alternatives, going back and forth between them. However, a look at the search paths suggests that this result may mainly be driven by an end-of-search effect, as almost everyone finishes her search by returning to the product she chose, which of course is satisfactory.

Similarly to the global search, fixating on an alternative that is already determined to be unsatisfactory is less likely in local search ($\bar{\phi}_{l3} = -.62$). Confirming the characterization of the local search, the probability of staying in the “local” area is about $\sim 90\%$ based on the results for $\bar{\phi}_{l0}$ to $\bar{\phi}_{l2}$. Refixating on the same product is most likely ($\bar{\phi}_{l1} = 1.14$), and moving to the right is somewhat more likely than moving to the left ($\bar{\phi}_{l2} = .04$). The remaining parameters have to be interpreted relative to the normalized brand parameters and with the previous ones in mind. The negative signs of $\bar{\phi}_{l4,Flavor}$ and $\bar{\phi}_{l4,Price}$ are probably due to the relative differences in size of the AOI relative to the remaining package. However, when combined with $\bar{\phi}_{l1}$, the sum is still positive, suggesting that targeted information search within the same product is still more likely than moving to the next product. Finally, we find support for strategic search by attribute only for search by price ($\bar{\phi}_{l5,Price} = 1.14$), but not by flavor ($\bar{\phi}_{l4,Flavor} = -.38$). This is probably caused by the considerable clutter on the shelves, which makes it difficult to move from one flavor AOI to another flavor AOI, in particular since the location of the flavor information on the package differs by brand. For price, this is a lot easier since all prices are in one line at the bottom of each shelf. Notice that the effect for search by price is of the same magnitude as the effect for refixating on the same product; thus, if the last fixation was on a price tag, a consumer is equally likely to stay on the same product or to move to the neighboring price tag. However, there is considerable heterogeneity across participants in how much strong this systematic search effect is ($\sigma_{l4,Price} = .33$), suggesting that some people may be more likely to use search by attributes than others.

As the final part for the search model, let's move to the transition probabilities. In general, transitioning to the local state is a lot more likely than transitioning to the global state ($\bar{\lambda}_{l0} = 1.54$). This is reflected in the finding that almost 80% of all fixations are estimated to occur in the local state. However, there seems to be no consistent impact of the previous state on those transition probabilities ($\bar{\lambda}_{l1} = -.08$). The high level of heterogeneity across participants ($\varsigma_{l1} = .88$) suggest that there may be inter-individual differences; looking at individual-level estimates, we find that 14.1% of participants are more likely to stay in the global state if the last fixation was also in the global state (relative to moving to the global state if the last fixation was in the local state), while 12.5% are more likely to move to the local state.

Similarly, participants are equally likely to quit their search after the local state as after the global state ($\bar{\lambda}_{t1} = .11$), with the initial stopping probability not surprisingly being extremely low ($\bar{\lambda}_{t0} = -5.17$). More interestingly, however, is a look at how the stopping probability changes over time. As one might expect, people become more likely to stop their search the longer they have already searched ($\bar{\lambda}_{t2} = 1.68$); this effect holds true for 90.6% of the participants, where the others do not seem to be affected by the length of their search. In contrast, we find support for the increase in stopping probability implied by the satisficing choice rule for *all* participants, i.e. having found the first satisfactory alternative significantly increases the stopping probability ($\bar{\lambda}_{t3} = 2.34$). To understand the relative magnitudes of these two effects, recall that the number of fixations is relative to the average number of fixations for a given choice set. Thus, the impact of finding the first satisfactory alternative on the stopping probability is 1.4 times larger than the effect of having searched the average search length (or alternatively, about 50 times larger than the impact of one additional fixation).

In the independent model, the satisficing stopping rule is not part of the model since λ_{t3i} is set to zero. We find that in this case the initial probability to stop searching immediately is significantly higher ($\bar{\lambda}_{t0} = -3.78$), i.e., the intercept picks up some of the effect that is now missing from the model. More interestingly though, we also find that the effect of the length of the search increases significantly (i.e., no overlap of the 95% highest density intervals) to $\bar{\lambda}_{t2} = 2.47$, also picking up some of the missing effect. The fact that the probability of having

Table 4: Posterior Means (and Standard Deviations) for the Choice Parameters

	<u>Brand</u>		<u>Flavor</u>		<u>Price</u>		
$\hat{\gamma}_{Fantastic}$.79	(.05)	$\hat{\gamma}_{Beef}$.46 (.06)	$\vec{p}_{4.00}$.03	(.02)
$\hat{\gamma}_{Indomie}$.73	(.05)	$\hat{\gamma}_{Cheese}$.60 (.06)	$\vec{p}_{4.75}$.04	(.02)
$\hat{\gamma}_{Koka}$.58	(.06)	$\hat{\gamma}_{Chicken}$.88 (.04)	$\vec{p}_{5.50}$.23	(.05)
$\hat{\gamma}_{Maggi}$.98	(.01)	$\hat{\gamma}_{Curry}$.65 (.06)	$\vec{p}_{6.25}$.21	(.05)
			$\hat{\gamma}_{Lobster}$.35 (.06)	$\vec{p}_{7.25}$.48	(.06)
			$\hat{\gamma}_{Mushroom}$.46 (.07)			
			$\hat{\gamma}_{OnionChicken}$.84 (.05)			
			$\hat{\gamma}_{Shrimp}$.44 (.06)			
			$\hat{\gamma}_{Tomato}$.45 (.06)			
			$\hat{\gamma}_{Vegetable}$.45 (.06)			

found at least one satisfactory alternative increases with the length of the search is trivially true, as by definition there is no satisfactory alternative at the beginning of the search and almost always at least one at the end of the search. Taken together with the results from the full model and the independent model, this essentially gives the standard three equations used for testing for mediation (Sobel 1982). The results can then be interpreted as evidence that the effect of the length of search is partially mediated by the indicator for having found at least one satisfactory alternative.

Taken together, these findings lend strong support to the hypothesis that consumers do in fact use the stopping rule implied by the satisficing choice rule.

6.2 Choice

Table 4 presents the posterior means and standard deviations for the population level hierarchies for the choice model parameters.

By and large, these estimates conform roughly to the choice shares presented in Tables 1 and 2. *Every* participant chose Maggi at least once, which is reflected in the high probability of acceptability ($\hat{\gamma}_{Maggi} = .98$). At first glance, it might be surprising that the population level probability of the highest price being available is almost 50%. However, a closer look

at the data reveals that in fact 48% of the participants chose a product priced at QR 7.00 at least once (implying that that price is acceptable to them), so the estimate is perfectly on target. Keeping in mind that even this highest price is only \$1.90 for a five pack of noodles, this is not all that surprising. In contrast, only one person never chose a product that cost more than QR 4.00.

To further test the face validity of our results, we correlate the individual-level results with the explicit measures of brand and flavor preference collected in the questionnaire. Across all participants, the correlation is .47 for brands and .57 for flavors. These correlations are strong considering the numerous ties in the explicit measures due to using a five point Likert scale and, more importantly, the numerous ties in the model estimates due to its deterministic nature (if a person chose several different flavors, they all have a probability of being acceptable of 1). For the individual-level correlations based on only four and ten values, respectively (to avoid scale issues across participants), the mean correlation is .60 for brands and .61 for flavors.⁸

Note that the model is absolutely deterministic in one direction: If someone chose a flavor/brand at least once, that flavor/brand has to be acceptable for that person. While this feature may seem odd when thinking in compensatory terms, it perfectly makes sense if one truly believes that a person uses a non-compensatory rule. If the flavor/brand were not acceptable, a product with that flavor/brand could have never been chosen.⁹

However, the reverse is not true. One might think that the model should always estimate that a flavor/brand never chosen was unacceptable to that particular person. Yet, that is not the case. There are two reasons for that: (1) If a person acquires very little information before making a choice, he may rarely have encountered a certain flavor, if at all. In that case, there is simply little information on the respective parameter, making the estimation largely reliant on the hierarchy. (2) Including the status of an alternative into the search cost can provide additional information about whether a certain flavor/brand was acceptable or not, even if it was never chosen. Say a person never chose mushroom flavor, but whenever

⁸Individual level correlations could not be calculated for 27 individuals for brands and for one individual for flavors due to no variation in the explicit and/or estimated preference measures.

⁹See appendix A.2 for brief description and results of a probabilistic version of the model.

she sees mushroom flavor, she also gathers the corresponding price information rather than moving on to the next product. In that case, it should be very likely that she does actually find mushroom an acceptable flavor. To understand this dynamic, we take a closer look at the individual-level parameters of acceptability for brands and flavors that were never chosen.

For the full model, we find nonetheless that for 68% of these cases, the probability that a non-chosen brand or flavor is acceptable is below 5%. However, the remaining 32% have considerable variation, with a mean probability of acceptability of 34% and even 2% of cases for which this probability is over 90%.

For the independent model, reason (2) mentioned above does not apply anymore. Thus, by comparing the two models, we can analyze how much of this variation is due to not very informative data and how much of it is due to the joint modeling of search and choice. Looking at the hierarchy parameters, we find that all flavors are estimated to be more likely to be acceptable than in the full model, on average by 7.5%. Since the independent model is also deterministic for flavors chosen at least once, this increase in acceptability must be caused by participants who never chose the respective flavor. A look at the individual-level estimates confirms this insight. For the non-deterministic cases, only 20% have a probability of the non-chosen brands and flavors being acceptable of less than 5% (down from 68%). Thus, using a status-dependent search helps overcome the potential problem of sparse data and draw the individual-level estimates away from the hierarchy.

As a final test for face validity, we identify three participants as vegetarians (defined by never choosing a non-vegetarian flavor and giving the lowest possible explicit rating to all non-vegetarian flavors). Naturally, the model should also be able to identify these individuals. Results are promising, yet lend further insight into reason (1) for non-zero acceptability probabilities given above. For all but one flavor for one person (out of five flavors times three people), the probabilities that non-vegetarian flavors are acceptable are very low. For the exception, this probability is 54%. Despite being the 8th lowest value for chicken across participants, this is still higher than one would like. Inspection of the search paths for this particular participant explains why: In six of the choice sets, s/he never even saw an option

with chicken flavor at all! Thus, there is not enough information in the data to draw the estimate further away from the very high population value in the hierarchy ($\hat{\gamma}_{Chicken} = .88$). In order to overcome the population value for individuals such as these, one either needs more data to allow preferences to be more fully observed or one could explicitly model potential preference structures, e.g., by adding an extra layer to estimate whether a certain person is vegetarian or not and including a parameter for whether a flavor is a vegetarian option or not.

The situation is exacerbated in the independent model since we also miss the additional information from the search. While the non-vegetarian flavor acceptability probabilities for the vegetarians are consistently below the respective hierarchy levels, they are far from identifying vegetarians as such. On average, non-vegetarian flavors are estimated to be acceptable for vegetarians with a probability of 47%, with one estimate even being over 90%. Once again, this highlights the importance of modeling search and choice jointly.

Finally, to gain further insight into whether consumers may or may not be using the proposed satisficing choice rule, we analyze the number of satisfactory options a person has found before stopping his search. On average, people have 1.75 satisfactory options to choose from at the end of their search. On an individual level, more than 70% of participants average less than two satisfactory options across choice sets before terminating their search. Once again this suggests that having found one satisfactory alternative is sufficient for many people to stop their search very soon after, lending further support to the hypothesis that they follow a satisficing choice rule. On the other hand, though, 8% of the participants have on average more than three satisfactory options before making their final choice, suggesting that a satisficing choice model may not be appropriate for them.

6.3 Holdout Prediction

In order to perform a holdout prediction task, we re-estimate the model using only twelve of the 15 choice sets and use the remaining three choice sets for prediction. Since the model estimates consumer preferences at the attribute level and choice probabilities are not compensatory across attributes, it is instructive to analyze predictive performance on the

attribute level.¹⁰ Therefore, we first discuss the results on the attribute level as well as aggregated to the product level, and then move to choice predictions and comparisons to two multinomial logit models.

6.3.1 Attribute- and Product-Level Predictions

Since the model provides probabilities of acceptability for each level of each attribute, we can check how well the model fits the holdout choices on an attribute level. The holdout choices conform extremely well with the model results. Recall that we estimate the posterior probability that an attribute is acceptable. We define the individual-level acceptable sets for a given attribute as those that are acceptable with probability of at least 95%. We find that 88% of the flavors, 97% of the brands, and 98% of the prices chosen in the holdout choices are within the respective acceptable sets.¹¹ Of course, the model strongly benefits from its deterministic nature, i.e., if a flavor chosen in the holdout choices was chosen by the same individual in one of the estimation choices, the probability of it being acceptable is necessarily 1. The somewhat lower hit rate for flavors is then mainly caused by the greater number of flavors to choose from and the resulting higher probability that a flavor chosen in the holdout choices may not have been chosen in the estimation choices.

Using the attribute-level results as well as the data on which pieces of information participants looked at for the holdout choices, we can calculate the product-level probabilities of each product for being satisfactory, unsatisfactory, or undetermined for each participant. Examining the products chosen in the holdout choices, we find that more than 75% of the choices have a probability above 95% of being satisfactory for the respective participant. The chosen product has the highest probability of being satisfactory in 82.2% of all cases; however, in more than half (53.3%) of those cases it is tied for first place with at least one more product. Once again, this is due to the fairly deterministic nature of the model.

¹⁰For a compensatory model, it is less obvious how to evaluate the predictive ability of the part-worth utilities since choices are a compensatory function of all attributes.

¹¹We exclude the no-choice instances that occur in the holdout choices for this subsection as well as the following one since the analyses are not applicable to them.

6.3.2 Choice Predictions

Using the choice rule proposed in the model, we also generate simulated choice probabilities for the holdout choices. Choice probabilities depend on the number of satisfactory and undetermined options at the time of decision as well as the trembling hand parameter. Using the product-level probabilities calculated in the previous section, we simulate the outcome of each holdout choice for each participant 100,000 times, calculate the resulting choice probability, and average across simulations.

The resulting choice predictions are distinctly multi-modal, which is not surprising given the choice rule. If a chosen product was likely to be the only satisfactory one in a choice set, its choice prediction will be close to 100%. If there are two products with a high likelihood of being satisfactory, the choice predictions will be close to 50%, and so on. On average, the predicted choice probability for the chosen option is 48.0%.¹²

In order to evaluate the predictive performance of the proposed model, we compare it to two different specifications of the multinomial logit model. In the model called *MNLall*, we employ the traditional assumption that a decision maker was aware of all information when making his choice, whereas in the model called *MNLind* we construct individual information sets for each participant based on what s/he actually looked at.¹³ The *MNLind* model should certainly do better than the *MNLall* model, but the comparison between these two models gives additional insight into how much prediction is improved in the traditional framework due to knowing the actual information sets of the decision maker.

¹²Notice that these simulations do not include any information on the search sequence (only on what information was acquired). Thus, if several products have a high probability of being satisfactory, in all likelihood the choice set will include all of these products in the simulations; however, as discussed above, choice sets are usually very small at the end of the search. Incorporating search information into the simulations may improve the choice predictions by adding more information on which alternatives actually were likely to have been satisfactory. This problem is particularly strong for the no-choice decisions in the holdout set. In the majority of no-choice decisions, the choice set actually included no satisfactory option (61% with average choice set of less than .1 satisfactory options when choosing no-choice). However, in our simulations this is not necessarily the case. Since that reduces the no-choice probability to the trembling hand probability, the holdout prediction for all but one of the cases in which no-choice was chosen is smaller than 10%.

¹³Products that were never fixated on are not included in those sets. If only one or two pieces of information were missing for a certain product, equal probability was given to all possible levels, i.e., we assume the decision maker is risk-neutral and uses the expected value to compute the product's utility.

The predicted choice probability for the chosen option is on average 44.7% for the *MNLall* model and 57.7% for the *MNLind* model, indicating a significant increase thanks to the additional information. Moreover, the *MNLall* model does better in terms of the average choice probability than the proposed model. However, the absolute choice probabilities may not be the best way to compare these models, given the multi-modal nature of the predictions from the proposed model.

One alternative is to look at hit rates, i.e., the probability that the chosen option has the highest predicted choice share. The resulting hit rates are 50.0% for the *MNLall* model, 58.9% for the *MNLind* model, and a staggering 78.6% for the proposed model. However, the extremely high hit rate for the proposed model does not take into account that in many of these correct predictions, the chosen product is tied with one or more other products for the highest choice probability (as would be expected given the ties in the probabilities of being satisfactory reported in the previous section). So while it is a very encouraging result that the model picks the chosen option to be among the top choices in almost 80% of the cases, looking only at these hit rates is not a fair comparison for the multinomial logit models.

The probably fairest comparison is the one also most relevant to retailers. A retailer likely not only cares about which product is most likely to be chosen, but would like to know all products that have a non-negligible choice probability. We therefore define the success rate of the prediction as the percentage of holdout choices that have a choice probability over a certain threshold probability. Notice that this definition inherently “punishes” the proposed model for having multiple-way ties, since having several options tied for the highest choice probability lowers the maximum absolute choice probability. Table 5 summarizes the results for different values of the threshold probability.

For (unrealistically) large threshold probabilities, the *MNLall* model does best, as the proposed model is severely punished for having ties. Yet even at the still sizable threshold of 25% the proposed model has a higher success rate than either of the multinomial logit models. 25% is certainly a large enough choice probability for a retailer to care about. The lower the threshold, the larger the advantage of the proposed model. While the satisficing model benefits from having larger mean set sizes (i.e., the number of options with choice proba-

Table 5: Success Rate (and Mean Set Size) for Different Thresholds

Threshold	<i>MNLall</i>	<i>MNLind</i>	Satisficing
.3	55.2% (1.125)	66.1% (1.161)	65.6% (1.260)
.25	59.4% (1.281)	66.7% (1.219)	68.2% (1.354)
.2	62.5% (1.438)	67.7% (1.307)	76.0% (1.693)
.15	67.2% (1.641)	70.8% (1.417)	81.25% (2.042)
.1	72.4% (1.948)	74.5% (1.536)	84.4% (2.422)

bilities above the threshold), the sets are still fairly small and practical from a managerial perspective.

The lower average set sizes for the multinomial logit models suggest that those models generate more extreme predictions. This is confirmed by looking at the choice probability of the option a person is most likely to choose according to the *MNLind* model. On average, these probabilities are 83%, and the median is even 90%. In turn, this of course leaves only fairly small choice probabilities for the remaining options, i.e., the model clearly singles out the number one choice. As evidenced by the hit rates reported earlier, these predictions are oftentimes correct. However, when the model prediction is incorrect, the choice probability of the chosen option is necessarily small, leading to the smaller average set sizes and the lower success rates.

In comparison, the mean of the highest choice probability per choice set for the satisficing model is only 55% (median = 48%). By only estimating which alternatives are satisfactory (i.e., good enough) and not further discriminating between those, the satisficing model results in lower predicted choice probabilities for the “best” option but allows itself to assign non-trivial choice probabilities to all potentially satisfactory options. Recall that the chosen alternative nonetheless has the highest choice probability (possibly tied) in 78.6% of the cases, i.e., the proposed model is very capable of identifying the correct choice options. Yet, the higher success rates for the satisficing model also suggest that not over-committing to just one choice prediction may be the more viable approach for predicting consumer choice.

7 Discussion

The proposed model continues the line of research by Gilbride and Allenby (2004); Jedidi and Kohli (2005). This line of research truly brings a paradigm shift to the empirical choice model literature in marketing, a shift away from compensatory utility maximizing and towards a quest for more realistic models of consumer choice. Most models in this new line of research employ a two-stage approach in which the simple heuristic is used to form a consideration set in the first stage, followed by a compensatory utility maximizing choice in the second stage. In contrast, the proposed model does not rely on compensatory tradeoffs at all. This is possible thanks to a search stopping rule based on Simon’s idea of a satisficing decision maker (Simon 1955). In a satisficing choice rule, the sequence in which products are evaluated is essential. We therefore collect choice and eye-tracking data in a visual conjoint experiment and jointly model search and choice.

The results lend significant support to the proposed model. Most importantly, the stopping rule implied by the satisficing rule is strongly supported by the parameter estimates. In addition, the distinction between satisfactory and unsatisfactory products is meaningful in explaining the search pattern, too. We also show that the joint model of search and choice informs the parameters of the choice model much better than the independent model.

The model performs extremely well in a holdout prediction task. The chosen option has the highest choice probability in almost 80% of the cases. Moreover, the proposed model is significantly more successful than a multinomial logit model in assigning non-trivial choice probabilities to the chosen option for threshold values that are relevant for retailers (say, 10-20%).

It has long been accepted that consumers do not really calculate the compensatory utilities implied by the standard models. Our results show that it is possible to estimate choice models that conform more closely to the actual decision making process - and that it may be worthwhile to do so! We therefore fully agree with Netzer et al. (2008) that it is time to improve what they call the “ecological fit” of the choice models to the respective task.

Of course we do not intend to imply that all consumers always follow a satisficing decision

rule. Heterogeneity across people in their tendency to use simple choice heuristics (often imprecisely called “satisficing”) vs. maximizing decision rules have been well documented (e.g., Schwartz et al. 2002). Moreover, the same person is likely to employ different choice rules when buying instant noodles vs. a car, for instance. And even for the same task, choice rules have been found to vary depending on time pressure, fatigue, etc. (e.g., Swait and Adamowicz 2001). Future research needs to address how to incorporate these issues into empirical choice models.

Given this heterogeneity in potential choice rules, we agree that a satisficing choice model may not always be the appropriate model when analyzing consumer choices. However, it should not come as a surprise that for frequently purchased (at least for the subject pool) and fairly inexpensive goods like instant noodles consumers employ simpler choice rules like the satisficing rule estimated in this paper. And if they do, our models should reflect that. Or so Simon says.

References

- Bettman, J. R., Johnson, E. J., and Payne, J. W. (1991). Consumer decision making. In Robertson, T. and Kassirjian, H., editors, *Handbook of Consumer Behavior*, chapter 2, pages 50–84. New York: Prentice Hall.
- Casella, G. and George, E. I. (1992). Explaining the Gibbs sampler. *American Statistician*, 49(4):327–335.
- Chandon, P., Hutchinson, J. W., Bradlow, E. T., and Young, S. H. (2009). Does in-store marketing work? Effects of the number and position of shelf facings on brand attention and evaluation at the point of purchase. *Journal of Marketing*, 73:1–17.
- Coombs, C. H. (1951). Mathematical models in psychological scaling. *Journal of the American Statistical Association*, 46(256):480–489.
- Dawes, R. M. (1964). Social selection based on multidimensional criteria. *The Journal of Abnormal and Social Psychology*, 68(1):104–109.
- Elrod, T., Johnson, R. D., and White, J. (2004). A new integrated model of noncompensatory and compensatory decision strategies. *Organizational Behavior and Human Decision Processes*, 95:1–19.
- Gelfand, A. E. and Smith, A. F. M. (1990). Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85:398–409.
- Gigerenzer, G. and Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In Gigerenzer, G., Todd, P. M., and ABC Research Group, editors, *Simple Heuristics That Make Us Smart*, chapter 1, pages 3–36. New York, NY, US: Oxford University Press.
- Gilbride, T. J. and Allenby, G. M. (2004). A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science*, 23(3):391–406.
- Gilbride, T. J. and Allenby, G. M. (2006). Estimating heterogeneous EBA and economic screening rule choice models. *Marketing Science*, 25(5):494–509.
- Guadagni, P. M. and Little, J. D. C. (1983). A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2(3):203–238.
- Gupta, S. (1988). Impact of sales promotions on when, what, and how much to buy. *Journal of Marketing Research*, 25(4):342–355.
- Hauser, J. R. and Wernerfelt, B. (1990). An evaluation cost model of consideration sets. *The Journal of Consumer Research*, 16(4):393–408.
- Heidelberger, P. and Welch, P. (1983). Simulation run length control in the presence of an initial transient. *Operations Research*, 31:1109–1144.
- Jedidi, K. and Kohli, R. (2005). Probabilistic subset-conjunctive models for heterogeneous consumers. *Journal of Marketing Research*, 42:483–494.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.
- Kamakura, W. A. and Russell, G. J. (1989). A probabilistic choice model for market segmentation and elasticity structure. *Journal of Marketing Research*, 26(4):379–390.
- Kohli, R. and Jedidi, K. (2007). Representation and inference of lexicographic preference models and their variants. *Marketing Science*, 26(3):380–399.

- Liechty, J., Pieters, R., and Wedel, M. (2003). Global and local covert visual attention: Evidence from a bayesian hidden markov model. *Psychometrika*, 68(4):519–541.
- Loomes, G., Starmer, C., and Sugden, R. (1991). Observing violations of transitivity by experimental methods. *Econometrica*, 59(2):425–439.
- Mehta, N., Rajiv, S., and Srinivasan, K. (2003). Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Science*, 22(1):58–84.
- Netzer, O., Toubia, O., Bradlow, E. T., Dahan, E., Evgeniou, T., Feinberg, F. M., Feit, E. M., Hui, S. K., Johnson, J., Liechty, J. C., Orlin, J. B., and Rao, V. R. (2008). Beyond conjoint analysis: Advances in preference measurement. *Marketing Letters*, 19:337–354.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press, Cambridge, MA.
- Pieters, R. and Warlop, L. (1999). Visual attention during brand choice: The impact of time pressure and task motivation. *International Journal of Research in Marketing*, 16:1–16.
- Pieters, R., Warlop, L., and Wedel, M. (2002). Breaking through the clutter: Benefits of advertising originality and familiarity for brand attention and memory. *Management Science*, 48(6):765–781.
- Roberts, J. H. and Lattin, J. M. (1991). Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28(4):429–440.
- Russo, E. J. and Leclerc, F. (1994). An eye-fixation analysis of choice processes for consumer nondurables. *Journal of Consumer Research*, 21:274–290.
- Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., and Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, 83(5):1178–1197.
- Shi, S. W., Wedel, M., and Pieters, R. (2010). A markov cascade analysis of information processes: An application to comparison websites. unpublished manuscript.
- Shugan, S. M. (1980). The cost of thinking. *The Journal of Consumer Research*, 7:99–111.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.
- Smith, B. J. (2007). boa: An R package for MCMC output convergence assessment and posterior inference. *Journal of Statistical Software*, 21(11):1–37.
- Sobel, M. E. (1982). Asymptotic confidence intervals for indirect effects in structural equation models. In Leinhardt, S., editor, *Sociological Methodology*, pages 290–312. Washington, DC: American Sociological Association.
- Sun, B. (2005). Promotion effect on endogenous consumption. *Marketing Science*, 24(3):430–443.
- Swait, J. (2001). A non-compensatory choice model incorporating attribute cutoffs. *Transportation Research Part B*, 35:903–928.
- Swait, J. and Adamowicz, W. (2001). The influence of task complexity on consumer choice: A latent class model of decision strategy switching. *Journal of Consumer Research*, 28:135–148.
- Teixeira, T. S., Wedel, M., and Pieters, R. (2010). Moment-to-moment optimal branding in TV commercials: Preventing avoidance by pulsing. *Marketing Science*, 29(5):783–804.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4):281–299.

- van der Lans, R., Pieters, R., and Wedel, M. (2008a). Competitive brand salience. *Marketing Science*, 27(5):922–931.
- van der Lans, R., Pieters, R., and Wedel, M. (2008b). Eye-movement analysis of search effectiveness. *Journal of the American Statistical Association*, 103(482):452–461.
- von Neumann, J. and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition.
- Wedel, M. and Pieters, R. (2008). A review of eye-tracking research in marketing. *Review of Marketing Research*, 4:123–147.

Appendix

A.1 Priors

To complete the hierarchical Bayesian setup, a set of priors is needed. We choose largely uninformative priors, as shown in Table 6. We scale the prior for $\hat{\tau}_2$ to be 100 times the prior for $\hat{\tau}_1$ to reflect the idea that the trembling hand probability should be fairly small. Nonetheless, the priors are wide enough to allow for a wide spectrum of Beta distributions on the trembling hand probabilities.

Table 6: Priors

<u>Choice</u>	<u>Search</u>
$\hat{\gamma}_a \sim \text{Beta}(1, 1)$	$\bar{\phi} \sim \text{Normal}(0, 100)$
$\bar{p} \sim \text{Dirichlet}(1, 1, 1, 1, 1)$	$\sigma \sim \text{Inverse Gamma}(.5, .5)$
$\hat{\tau}_1 \sim \text{Gamma}(2, 1)$	$\bar{\lambda} \sim \text{Normal}(0, 100)$
$\hat{\tau}_2 \sim 100 \cdot \text{Gamma}(2, 1)$	$\varsigma \sim \text{Inverse Gamma}(.5, .5)$

A.2 Probabilistic Satisficing Model

We also estimated a probabilistic version of the model based on Jedidi and Kohli (2005). In the probabilistic version, each attribute level is acceptable or unacceptable only with a certain probability, rather than with certainty. This makes sense if one assumes that a consumer decides on the spot whether a given price, flavor, or brand is acceptable each time it is encountered, allowing for some decision uncertainty or error. In contrast, the deterministic model assumes that these decisions have already been made *a priori* and therefore stay the same across all choices.

It is easiest to represent the changes to the choice part of the model by replacing equations 1 and 2 with updating rules for whether a product is satisfactory or not. All products are undetermined at the start, i.e. $D_{x0} = 0$ for all x . Once a product has been determined to be either satisfactory or unsatisfactory, that status cannot change anymore (i.e., the updating rules only apply if the product is still undetermined). By fixating on a given attribute of a product (a_x^*), that product has a chance of becoming unsatisfactory if the attribute is judged to be unacceptable. Only after fixating on the third attribute can a product become

satisfactory if all three attributes are acceptable. This is equivalent to saying that the product has not been judged to be unsatisfactory (implying the first two attributes were acceptable) and that the third attribute is acceptable. We thus have:

$$Pr(S_{xf} = 1 | D_{x,f-1} = 0) = Pr(a_x^* \in \Gamma_A^*) \cdot \prod_A \mathcal{I}_{a_x \in I_{jf}}$$

and

$$Pr(U_{xf} | D_{x,f-1} = 0) = Pr(a_x^* \notin \Gamma_A^*)$$

Notice that while the interpretation of the model is still non-compensatory (i.e., if one attribute is judged to be unacceptable, the product is necessarily unsatisfactory), its mathematical form now becomes compensatory since the probability of a product being un/satisfactory is the product of the attribute-level probabilities. Thus, changing the flavor to a flavor that is less likely to be acceptable can be offset by changing the brand to a brand that is more likely to be acceptable.

In addition, the hierarchies given in equations 3 and 5 need to be adjusted. For brands and flavors, the hierarchy becomes a Beta distribution. For price, we impose that the acceptance probability of a price can at most be as much as the acceptance probability of the next lower price. We therefore use a Beta hierarchy for the acceptance probability of the lowest price, and four more independent Beta distributions for the ratio of the acceptance probability of price level x relative to the acceptance probability of price level $x - 1$.

Not surprisingly, the main differences in the result are found in the choice part of the model. It is noteworthy, though, that the search parameters corresponding to the status of the alternative ($\bar{\phi}_{g1}$, $\bar{\phi}_{g1}$, and $\bar{\phi}_{l3}$) become larger in absolute value. This is most likely the result of the fact that participants tend to have even fewer satisfactory alternatives in their choice set at the end of the search (since not every flavor, brand, and price they ever chose automatically has to always be acceptable). While just over 70% of the participants had on average less than two satisfactory options per choice set in the deterministic version, all but one participant (98.4%) now fall in that range. In fact, more than 90% of the participants have an average of 1.5 satisfactory alternatives or less when making their choice!

Table 7: Posterior Means (and Standard Deviations) of the Means of the Hierarchy Distributions for Acceptability

	<u>Brand</u>			<u>Flavor</u>			<u>Price</u>	
Fantastic	.56	(.05)	Beef	.40	(.06)	$\hat{p}_{4.00}$.95	(.01)
Indomie	.55	(.05)	Cheese	.50	(.06)			
Koka	.37	(.05)	Chicken	.87	(.04)	$\frac{\hat{p}_{4.75}}{\hat{p}_{4.00}}$.80	(.04)
Maggi	.80	(.04)	Curry	.60	(.06)			
			Lobster	.30	(.05)	$\frac{\hat{p}_{5.50}}{\hat{p}_{4.75}}$.74	(.05)
			Mushroom	.38	(.06)			
			Onion Chicken	.74	(.05)	$\frac{\hat{p}_{6.25}}{\hat{p}_{5.50}}$.75	(.06)
			Shrimp	.37	(.05)			
			Tomato	.36	(.05)	$\frac{\hat{p}_{7.00}}{\hat{p}_{6.25}}$.58	(.07)
			Vegetable	.40	(.06)			

Table 7 presents the posterior means of the means of the Beta hierarchies for acceptability. First note that both brands and flavors have a lower average acceptance probability than in the deterministic model. Once again, this is due to the fact that in the deterministic model having chosen a flavor once means the flavor is acceptable with certainty, whereas in the probabilistic model the probability may only be slightly higher than if never chosen. For the price column, the interpretation is very different due to the differing structure of the hierarchy. To illustrate the results for prices, we give the medians of the individual-level parameters. The lowest price of QR 4.00 has median acceptability probability of 95.9%, which is not surprising. The second level (QR 4.75) still has a median acceptability probability of 84.4%, but for one individual it already drops down to 5.7%. The remaining median acceptability levels are 70.2% for QR 5.50, 46.3% for QR 6.25, and 22.1% for QR 7.00.

These results for the attribute level are very reasonable, assuming people decide about the acceptability of a certain attribute level on the spot every time. Not surprisingly, the probabilities of a chosen product to be satisfactory are lower relative to the deterministic model due to the added uncertainty. Due to fewer ties thanks to the probabilistic nature, the chosen product now has the highest probability of being satisfactory in only 69% of the cases; however, in almost all (94.4%) of these cases it is not tied for first place.

Table 8: Success Rate (and Mean Set Size) for Different Thresholds

Threshold	.3	.25	.2	.15	.1
Success Rate	55.2%	63.0%	73.4%	79.7%	85.0%
Mean Set Size	(.802)	(1.021)	(1.307)	(1.568)	(2.391)

While these results are very reasonable, the added uncertainty in the probabilistic model hurts its performance in the holdout prediction. Not surprisingly, the probabilities of a chosen product to be satisfactory are lower relative to the deterministic model. Due to fewer ties thanks to the probabilistic nature, the chosen product now has the highest probability of being satisfactory in only 69% of the cases; however, in almost all (94.4%) of these cases it is not tied for first place.

Table 8 presents the success rate for the choice predictions for the same thresholds reported in Table 5. The probabilistic version has a lower success rate than the deterministic version except for very low threshold values. Its low mean set sizes for higher thresholds suggests that its compensatory mathematical nature also leads to more extreme predictions, similar to the multinomial logit model. These results suggest that using compensatory models for holdout prediction may in fact overfit the estimates to the estimation data set (at least for small estimation data sets).