Consideration Sets and Competitive Marketing*

Kfir Eliaz† and Ran Spiegler‡

August 31, 2008

Abstract

We study a market model in which competing firms use costly marketing devices to influence the set of alternatives which consumers perceive as relevant. Consumers in our model are boundedly rational in the sense that they have an imperfect perception of what is relevant to their decision problem. They apply well-defined preferences to a "consideration set", which is a function of the marketing devices employed by the firms. We examine the implications of this behavioral model in the context of a competitive market model, particularly on industry profits, vertical product differentiation, the use of marketing devices and consumers' conversion rates.

KEYWORDS: marketing, advertising, consideration sets, bounded rationality, limited attention, persuasion, irrelevant alternatives

1 Introduction

We present a model of competitive marketing based on the notion that consumers are boundedly rational and that firms use marketing tactics in an attempt to influence consumers' decision process. The standard model of consumer behavior assumes that the consumer applies well-defined preferences to a perfectly perceived set of available alternatives. We retain the assumption that consumers have stable preferences, but relax the assumption that they have a perfect perception of what is relevant for their consumption problem, thus allowing firms to manipulate that perception. Our aim is to

*This is a substantially revised version of a working paper that appeared under the same title in 2006. We thank Amitav Chakravarti, Eddie Dekel, Sergiu Hart, Barton Lipman, John Lynch, Martin Peitz, Ariel Rubinstein and numerous seminar participants for useful conversations. Spiegler acknowledges financial support from the ISF, Grant no. 0610114022.

†Department of Economics, Brown University, Providence, RI 02912. E-mail: Kfir_Eliaz@Brown.edu. URL: http://www.econ.brown.edu/fac/Kfir_Eliaz.

‡Department of Economics, University College London, Gower St., London WC1E 6BT, UK. E-mail: r.spiegler@ucl.ac.uk. URL: http://www.homepages.ucl.ac.uk/~uctprsp/

explore the market implications of this departure from the standard model, especially for the way firms deploy marketing strategies in competitive environments.

The cornerstone of our model is the observation that in the modern marketplace, consumers face an overwhelmingly large variety of products and therefore often use screening criteria (deliberate as well as unconscious) in order to reduce the number of "relevant" alternatives. As a result, consumers apply their preferences not to the set of objectively feasible alternatives, but to a potentially smaller set which they construct at an earlier stage of the decision process. Borrowing a term from the marketing literature, we refer to this set as the "*consideration set*". The basic idea underlying this term is that consumers may be unaware of some of the feasible products, and even when they become aware of a new product, they still need to be *persuaded to consider* it as a potential substitute to their currently consumed product.

Our model of consumer behavior attempts to capture the idea that consumers do not automatically perceive all available options as relevant for their consumption problem, and they resist considering new alternatives. The role of marketing is to overcome this resistance. Whether or not it succeeds depends on the competing products and how they are marketed. The framework we propose accommodates a variety of ways in which marketing influences the formation of consideration sets. Here are a few examples.

Argumentation by a salesperson. Think of a consumer who enters a car dealership with the intention to buy a new car. The consumer has a particular car model in mind, and as he inspects it on the display podium, he is approached by a salesperson who tries to convince him to consider a different car model, using arguments that point out similarities and differences between the two models.

Positioning. Economists have extensively studied the way firms strategically differentiate their products in the space of product attributes. However, marketing a product often involves locating it in the more amorphous space of images and consumer perceptions. This type of product differentiation is known in the marketing literature as "positioning". For instance, a yogurt with given objective characteristics can be marketed with an emphasis on hedonic features such as taste and texture, or on health-related features, real or imaginary. Two yogurt brands may be differentiated in terms of their positioning even when their objective characteristics are very similar. Although this type of product differentiation may be payoff-irrelevant, it can affect the consumer's decision whether to consider a new brand.

Packaging. Consumers' decision to add a new product to their consideration set can also be influenced by the way it is packaged. For instance, a consumer is more likely to notice a new brand of flavored water with added vitamins if its bottle is designed so that it looks like the brand the consumer regularly buys.[1]

Advertising content. An ad that highlights a drawback in a group of rival products may attract the attention of consumers who currently consume those products ("tired of hours of waiting for customer service?" "having trouble keeping track of the fees your credit company charges you?"). Similarly, an ad for a product that highlights one of its good features may give consumers of other products that lack this feature a sufficient reason to consider the advertised product. The effect of ad content in these situations can go beyond mere information transmission. In particular, an ad that points out a flaw in a product the consumer regularly buys hardly tells him something he does not already know, and it is not likely to modify his beliefs about the quality of alternative products. However, it may have the rhetorical effect of persuading him to look for a substitute.

Products as attention grabbers. When firms design their product lines and retailers decide which products to put on display, they may take into account the possibility that certain products can help draw consumers' attention to other products on offer. For instance, think of a consumer who wants to buy a new laptop computer. He initially considers a particular model x, possibly because it shares some features with his current machine. The consumer may then notice that a computer store offers a model y that is significantly cheaper or lighter than x. This gives the consumer a sufficient reason to consider y in addition to x. Upon closer inspection, the consumer realizes that he does not like y as much as he does x. However, since he is already inside the store, he may browse the other laptop computers on offer and find a model z that he ranks above both x and y. Thus, although few consumers may actually buy y, this model functions as a "door opener" that attracts consumers' serious attention to the other products offered by the store.[2]

Search engine optimization. The internet has given rise to new marketing devices with which sellers try to expand consumers' consideration sets. Think of a consumer who wishes to spend a weekend in a quiet place out in the countryside. To find such a place on the internet, the consumer needs to enter keywords in a search engine. However,

[1] Compare the brand "VitaminWater" by Glaceau (http://www.glaceau.com) with the competing brand "Antioxidant Water" by Snapple (http://www.snapple.com).

[2] A vivid example of this effect involves a soda company that issues a "limited holiday edition" including absurd flavors such as Christmas ham or latke - see http://www.jonessoda.com/files/limited_editions.php

there is a variety of keywords he can use: “country inn”, “lodge”, “bed and breakfast”, “cottage”, etc. The consumer’s choice of keyword is likely to be guided by the labels he encountered in past vacations. Different keywords will elicit different lists, and suppliers can manipulate the list by bidding for keywords in sponsored-links auctions, or by employing a variety of “search engine optimization” techniques.

We propose a simple model of consideration set formation and embed it in a market environment in which firms employ marketing techniques to manipulate consideration sets. There are two identical firms and a continuum of identical consumers. Each firm chooses a pair, a product x and a marketing strategy M, and incurs a fixed cost associated with its choice. Each consumer is randomly assigned to one of the firms. The consumer’s initial consideration set includes only the product offered by that firm. This is interpreted as the consumer’s status quo, or default product. Whether or not the consumer also considers the competing firm’s product will depend on a primitive of his decision procedure, called the *consideration function*. This is a characteristic function that determines whether a consumer who initially considers a product x accompanied by the marketing strategy M will also consider a product x' when it is accompanied by the marketing strategy M'. If the consumer ends up including both firms’ products in his consideration set, he chooses his most preferred product, according to a strict preference relation defined on the set of products. Preferences are stable and impervious to marketing.

We analyze symmetric Nash equilibria of the game played by the two firms, when the firms’ objective is to maximize market share minus fixed costs. We use this model to address the following questions. Does the bounded rationality of consumers - namely their resistance to considering new, potentially superior products - enable firms to earn profits in excess of what they would earn if consumers were rational? Or does market competition (which includes marketing) eliminate this potential source of exploitation? Do firms’ profits necessarily decrease as consumers become “more rational” in the sense of being more likely to consider new products? Finally, we address the question of “converting consumers”. What is the fraction of consumers who switch a supplier in equilibrium? What is the probability that a consumer will buy a new product conditional on having been persuaded to consider it?

For the sake of tractability, we address these questions under a few further restrictions on the market model. We show that under certain conditions, there exists a symmetric Nash equilibrium in which firms earn the same profits they would earn if consumers were rational. We identify families of consideration functions for which *every* symmetric Nash equilibrium satisfies this property. In addition, we provide a complete characterization of symmetric equilibria for certain special classes of consideration functions. A notable feature of these characterizations is that they describe *both* the vertical product differentiation in the market and the marketing strategies that firms use to promote their products.

These findings lead us to investigate the effect of "enhancing consumer rationality" on industry equilibrium profits. We begin with a particular consideration function that generates rational-consumer profits in all symmetric Nash equilibria. Keeping everything else equal, we modify this consideration function so as to make it "more rational", in the sense that the consideration set coincides with the feasible set in a larger set of consumption problems. We show that this change gives rise to symmetric equilibria with higher, "collusive" industry profits. This example demonstrates that industry profits need not go down as consumers become "more rational".

Regarding the issue of consumer conversion, we show that in any symmetric equilibrium with rational-consumer profits, marketing is "effective" in the following sense. Consumers who add a new product to their consideration set thanks to active marketing end up buying the new product, as long as their default is not the most preferred product. Thus, symmetric equilibria with rational-consumer profits exhibit a strong correlation between persuading a consumer to consider a new product and persuading him to buy it.

The main contribution of our paper is the introduction of a framework for modelling the "persuading to consider" role of marketing, and the demonstration that it can illuminate aspects of competitive marketing. Our framework is flexible; it can incorporate a large variety of marketing methods, and in particular it enables us to address aspects of advertising content that are typically hard to capture with standard models.

The paper proceeds as follows. Section 2 describes the consumer's choice procedure. Section 3 presents the market model in which we embed the consideration-sets procedure. In Section 4, which is the analytical heart of the paper, we characterize symmetric Nash equilibria in the market model. Section 5 contains a detailed discussion of the paper's relation to the choice theory and marketing literatures, as well as the economic literature on advertising. Section 6 concludes.

2 Consumer choice

Let X be a finite set of *products*. Let $\mathcal{M}$ be a finite set of *marketing devices*. An *extended alternative* is a pair $(x, M) \in X \times 2^{\mathcal{M}}$ - i.e., a product and a collection of marketing devices that accompany it. Consumers in this paper will face choice problems that involve *ordered pairs* of extended alternatives, $((x_1, M_1), (x_2, M_2)) \in S^2$, where $S \subseteq X \times 2^{\mathcal{M}}$. The ordering has significance for us, as we will interpret (x_1, M_1) as the consumer's status quo or default alternative. Given a pair $((x_1, M_1), (x_2, M_2))$, the feasible set of products is taken to be $\{x_1, x_2\}$.

Consumer choice follows a two-stage procedure, which is based on two primitives: a *linear ordering* $\succ$ over X, and a *consideration function* $\phi : S^2 \to \{0, 1\}$. In the first stage, the consumer constructs a *consideration set*, which can take two values: $\{x_1, x_2\}$ if $\phi(x_1, M_1, x_2, M_2) = 1$, or $\{x_1\}$ if $\phi(x_1, M_1, x_2, M_2) = 0$. In the former case, the consideration set coincides with the objectively feasible set. In the latter case, it consists of the default option alone and thus forms a strict subset of the objectively feasible set. In the second stage of the choice procedure, the consumer chooses the $\succ$-maximal product in the consideration set. We say that (y, N) *beats* (x, M) if $\phi(x, M, y, N) = 1$ and $y \succ x$.

We interpret the linear ordering $\succ$ as the consumer's "true" preferences over X. The preference ranking $x \succ y$ is a genuine reflection of the consumer's taste, which is applied after serious consideration or actual experience with the products. This taste is stable and impervious to marketing. Put differently, if the consumer always considered all feasible products, then his revealed choices of products would be rationalized by $\succ$, which is also the preference criterion we adopt for normative analysis. Throughout the paper, x^* denotes the $\succ$-maximal product in X.

The consideration function ϕ enriches our description of the consumer's psychology. In addition to his preferences over products, the consumer is characterized by his willingness (or ability) to consider x_2 as a potential substitute to x_1, and how this willingness depends on the way the two products are marketed. Indeed, personality psychologists often regard "*openness to experience*" as one of the basic traits that define an individual's personality (see Goldberg (1993)). The consideration function may be viewed as a representation of this trait: ϕ' represents a more "open" personality than ϕ if $\phi(x_1, M_1, x_2, M_2) = 1$ implies $\phi'(x_1, M_1, x_2, M_2) = 1$.

Let us demonstrate the scope of this model with several specifications of $\mathcal{M}$ and ϕ. Some of these examples will be revisited in later sections.

Example 2.1: Pointing out flaws in the default product. Let $X = \{0, 1\}^K$ and $\mathcal{M} = \{1, ..., K\}$. Suppose that ϕ is only a function of x_1 and M_2. Specifically, assume that $\phi(x_1, M_2) = 1$ if and only if the m-th component of x_1 is 0 for some $m \in M_2$. The interpretation is that a product is a collection of attributes. A marketing strategy draws attention to a set of attributes and persuades the consumer to consider x_2 if and only if x_1 lacks at least one of them.

Example 2.2: Spurious product categorization. Let $\mathcal{M}$ be a finite set of categories. Suppose that S consists of all pairs (x, M) in which M is a singleton. Suppose further that ϕ is only a function of M_1 and M_2. Specifically, $\phi(M_1, M_2) = 1$ if and only if $M_1 = M_2$. The interpretation is that each product can be categorized in an arbitrary way that is irrelevant for consumer welfare. However, the categorization is relevant for the consideration function: the consumer considers x_2 if and only if it is assigned to the same category as x_1. An example is the location of products in supermarket isles. A sunscreen lotion for babies can be categorized as a skincare product or as a baby product. Similarly, a granola bar can be categorized as health food or a snack. Although the product's categorization is payoff-irrelevant, it determines its supermarket isle location, and consequently the likelihood that a consumer who initially considers one product will also notice the other.

Example 2.3: Irrelevant alternatives as marketing devices. Suppose products can be ordered along a line such that X can be taken to be a subset of the real numbers. Let $\mathcal{M} = X$ and let $S \subset X \times 2^{\mathcal{M}}$ contain all extended alternatives (x, M) having the property that M is contained in the $\succ$-lower contour set of x. That is, a product is presented against the background of inferior, "irrelevant" alternatives. Assume that ϕ is only a function of x_1, x_2 and M_2. More specifically, letting $I(x)$ be some neighborhood of x, assume $\phi(x_1, x_2, M_2) = 1$ if and only if $y \in I(x_1)$ for some $y \in M_2 \cup \{x_2\}$. One interpretation is that ϕ captures a similarity judgment between products: the consumer is willing to consider x_2 if and only if this product, or one of its accompanying irrelevant alternatives, are similar to x_1.

Example 2.4: Advertising intensity. Let $\mathcal{M} = \{1, ..., K\}$. Assume that ϕ is only a function of M_1 and M_2, such that $\phi(M_1, M_2) = 1$ if and only if $|M_2| \geq |M_1|$. The cardinality of M represents advertising intensity. The consumer considers the new product x_2 if and only if it is advertised at least as intensively as the status quo product x_1.

Our choice model displays a status-quo bias. Given (x_1, M_1, x_2, M_2), the consumer chooses the default/status-quo product x_1 whenever $x_1 \succeq x_2$. However, the consumer may continue to choose x_1 even when $x_2 \succ x_1$, if it happens to be the case that $\phi(x_1, M_1, x_2, M_2) = 0$. This is a status-quo bias of a different kind than the one studied in the literature (see Masatlioglu and Ok (2005)), where the status-quo tends to be *preferred* to other alternatives. In our choice model, the bias in favor of the status-quo exists in an earlier stage of the decision process, in which the consumer constructs the set of alternatives he will later consider for choice. Thus, the alternative to the status-quo is at a disadvantage not because the consumer tends to find it inferior to the status quo, but because he does not always take it into serious consideration.[3]

We conclude this section by commenting on two main assumptions in our model. First, the model draws a distinction between the product x and the marketing devices M that promote it. In reality, the boundary between the two may be blurred. For example, is the packaging of a product a pure marketing strategy, or is it part of the product's description? Any application of the consideration-sets model involves a modeling judgment as to which aspects of the product are payoff-relevant and which are viewed as pure marketing.[4]

Second, we assume the consideration function ϕ is exogenous. It would be interesting to derive this function as a result of some optimization that is carried out by a consumer (possibly a boundedly rational consumer who takes into account cognitive constraints). Such an analysis is beyond the scope of this paper. Still, there is some justification for treating ϕ as exogenous in our framework. The consideration function captures basic principles of what draws an individual's attention away from his default alternative. One example is the principle of similarity: people are often more inclined to pay serious attention to things that are similar to what they are already familiar with. Likewise, people are more likely to notice a statement about some product flaw when they regularly consume a product that actually has that flaw. These are general principles that are not specific to one market or another. Whatever optimization lies behind them is not market-specific, but takes place on a much larger, "general equilibrium" or "evolutionary" scale, where the consideration function is designed to be optimal on average across a large variety of market situations. As long as we analyze an individual market, it makes sense to treat the consideration function as exogenous.

3 A market model

The heart of this paper is a market model that incorporates the choice procedure introduced in Section 2. Our market consists of two identical firms and a continuum of identical consumers. The firms play a symmetric simultaneous-move game. The strategy space is $S \subseteq X \times 2^{\mathcal{M}}$, which is assumed to be sufficiently rich in the sense

[3] There is also a formal difference between the two notions of status-quo bias, which we discuss in Section 6.

[4] For instance, when the products in question are laundry detergents, it makes sense to assume that packaging and advertising are payoff-irrelevant. This assumption is less obvious when the products in question are, say, perfumes.

that $(x, \emptyset) \in S$ for every $x \in X$ and $(x^*, M) \in S$ for every $M \subseteq \mathcal{M}$. One reason for restricting the set of strategies is that a particular marketing strategy M may be inherently infeasible for promoting a given product x. For instance, when marketing involves highlighting certain product features, it is natural to assume that a firm cannot highlight a feature that its product lacks. The assumption that firms have identical strategy spaces is interpretationally non-trivial, because it rules out firm-specific brand names as marketing devices.

Each consumer is initially assigned to one of the firms (where each firm receives half the population of consumers). The extended alternative chosen by this firm plays the role of the default in the consumer's choice procedure. Choosing a strategy $(x, M) \in S$ entails a fixed cost of $c(x, M) \in (0, \frac{1}{2})$ given by $c(x, M) = c_x + \sum_{m \in M} c_m$. All c_x and c_m are strictly positive. Assume that $x \succ y$ implies $c_x \geq c_y$, with strict inequality for $x = x^*$. Firms aim to maximize their market share minus fixed costs.

We focus on the case in which the consideration function ϕ is independent of M_1. This assumption restricts the role of marketing to persuading consumers to consider new products, such that the marketing devices that accompany the default product do not affect the consumer's consideration set. This assumption rules out situations in which marketing can be used to dissuade consumers from considering new products (as in Examples 2.2 and 2.4 in the previous section). In addition, we assume that ϕ is "separable" in the following sense: if $\phi(x_1, x_2, M_2) = 1$ and $\phi(x_1, x_2, M_2 \backslash \{m\}) = 0$, then $\phi(x_1, x_2', M_2') = 1$ whenever $m \in M_2'$. That is, the effect that an individual marketing device has on the consideration set is independent of the new product and the other marketing devices that accompany it.

The market model is thus fully characterized by the tuple $\langle S, c, \succ, \phi \rangle$. The latter three components induce a payoff function in the simultaneous-move game played by the two firms. Throughout the paper, we use σ to denote a mixed strategy (namely, a probability distribution over S), and $Supp(\sigma)$ to denote its support. We favor the population interpretation of symmetric mixed-strategy equilibrium: there is a "sea of firms" from which two are randomly selected to play the roles of a default and a contender. Finally, $\beta_\sigma(x) = \sum_M \sigma(x, M)$ is the probability that the product x is offered under σ.

An important benchmark for this model is the case of a rational consumer. This case is subsumed into our model by letting $\phi(x, M, y, N) = 1$ for all x, M, y, N. A consumer with such a consideration function always considers both x_1 and x_2, and therefore always chooses according to $\succ$. Under consumer rationality, each firm plays the pair $(x^*, \emptyset)$ in Nash equilibrium, and consequently earns a payoff of $\frac{1}{2} - c_{x^*}$. We refer to the latter as the *rational-consumer payoff*.

The following condition will be maintained throughout the paper.

Assumption (*): $\phi(x, M, x^*, \mathcal{M}) = 1$ for all x, M.

Assumption (*) means that when the new product is the most preferred one, and in addition, it is accompanied by the grand set of marketing devices, then it is guaranteed to be considered. Note that this is the costliest strategy in S. Together with the assumption that $c(x, M) < \frac{1}{2}$ for all (x, M), Assumption (*) implies that $(x^*, \emptyset)$ and $\frac{1}{2} - c_{x^*}$ are a max-min strategy and the max-min payoff (we already observed that these are the Nash equilibrium strategy and Nash equilibrium payoff, respectively, under the rational-consumer benchmark). When we turn to analyzing symmetric Nash equilibria in our model, an important question will be whether firms are able to attain equilibrium payoffs *above* the max-min level.

Assumption (*) has another implication, which will be repeatedly employed in the sequel: the most preferred product is offered with positive probability in any symmetric Nash equilibrium.

Lemma 1 *Let σ be a symmetric Nash equilibrium strategy. Then, $\beta_\sigma(x^*) > 0$.*

Proof. Assume the contrary. let y denote the $\succ$-minimal product for which $\beta_\sigma(\cdot) > 0$. The market share that any $(y, M) \in Supp(\sigma)$ generates in equilibrium is at most $\frac{1}{2}$. If a firm deviated to $(x^*, \mathcal{M})$, it would ensure a market share of one. By the assumption that $c(x, M) < \frac{1}{2}$ for all (x, M), this deviation is profitable. ■

Note that this result can be extended to more general environments, as it only depends on Assumption (*) and our assumption that $c(x, M) < \frac{1}{2}$ for all (x, M).

We conclude this section with a comment on the model's interpretation. We view the market model as a vehicle for an abstract exploration of some market implications of the consideration-sets choice model (in fact, our market model could also be used with other boundedly rational choice models). We expect it to deliver insights into the effects of marketing in competitive environments. However, the model is *not* meant to be a faithful description of some concrete set of industries. This is evident from two features of the model. First, we are silent as to whether firms should be viewed as manufacturers or retailers. Whether the former or the latter is a more reasonable interpretation will vary from one example to another. Second, we abstract from price setting, for the sake of analytic simplicity. As with spatial competition models, it is easier to begin by assuming that firms care only about market share and defer the incorporation of prices into the model.[5]

4 Equilibrium analysis

In this section we analyze the symmetric Nash equilibria in the market model. In Sub-Section 4.1, we study the relatively simple case in which the consideration function is independent of x_2. In Sub-Section 4.2, we turn to the more complicated case in which the consideration function is also allowed to depend on x_2, in addition to x_1 and M_2. Sub-Section 4.3 is devoted to a property of symmetric Nash equilibrium that concerns the correlation between marketing and product quality. We discuss our findings in Sub-Section 4.4.

4.1 Influencing consideration sets through pure marketing content

In this sub-section we assume that whether consumers consider a new product depends only on the default product and the marketing devices that accompany the new product. That is, ϕ is only a function of x_1 and M_2; specifically, $\phi(x_1, M_2) = 1$ if and only if there exists $m \in M_2$ satisfying $\phi(x_1, \{m\}) = 1$. Hence, active marketing is necessary for a new product to enter the consumer's consideration set. Henceforth, we will say that the marketing device m is *effective* against x whenever $\phi(x, \{m\}) = 1$. Let $X_\phi(m)$ denote the set of products against which m is effective, i.e.,

$$X_\phi(m) \equiv \{x \in X : \phi(x, \{m\}) = 1\} \tag{1}$$

Without loss of generality, we assume that $x^* \notin X_\phi(m)$ for all $m \in M$. This assumption is made purely for notational convenience. Note that Assumption (*) implies that $\{X_\phi(m)\}_{m \in \mathcal{M}}$ forms a *cover* of $X \setminus \{x^*\}$.

This restricted domain captures a variety of marketing situations. One interpretation of $\mathcal{M}$ is as a set of advertising messages. A message m may be that the target (default) product has a certain feature, such that $\phi(x_1, \{m\}) = 1$ whenever the message is correct. This captures situations in which advertising reminds the consumer of a problem with his current product, and this reminder whets his appetite for a change.

[5]The objective function *is* realistic in some situations, e.g. when firms are platforms such as broadcast TV channels or internet portals. An alternative scenario, suggested to us by Bruno Strulovici, is that consumers are endowed with a fixed budget and spend all of it on their chosen product.

This specification of the model also fits certain aspects of positioning. We can interpret $\mathcal{M}$ as a set of possible product images. Consumers respond only to images that firms attach to new products. In contrast, the image that is attached to the consumer's default product does not affect him because he is familiar with the actual product. Thus, $\phi(x_1, \{m\}) = 1$ whenever the image m that is attached to the new product persuades the consumer to consider it, given that he currently consumes x_1. For instance, m may create an impression that the new product is healthier or more glamorous than x_1, or conversely that it is quite similar to x_1, and that may be enough to convince the consumer to try the new product.

The following is a simple example that captures an advertising technology in the manner of Butters (1977): consumers become aware of a new product if and only if it is advertised (note that unlike our model, in Butters (1977) consumers are not initially attached to any firm: if no firm advertises, consumers stay out of the market).

Proposition 1 *Suppose that* M *consists of a single marketing device* m*. Then, there is a unique symmetric Nash equilibrium, given by:*

$$\begin{aligned}\sigma(y^*, \emptyset) &= 2c_m \\ \sigma(x^*, \{m\}) &= 2(c_{x^*} - c_{y^*}) \\ \sigma(x^*, \emptyset) &= 1 - 2(c_{x^*} - c_{y^*} + c_m)\end{aligned}$$

where y^* *is the* $\succ$*-minimal product in* X*.*

We omit the proof of this result since it is a special case of Proposition 4, which is proven below.

The above equilibrium has several noteworthy properties:

1. The equilibrium strategy is mixed and consumers end up buying an inferior product with positive probability.

2. Firms advertise with positive probability.

3. Although the equilibrium outcome departs from the rational-consumer benchmark, firms earn the rational-consumer (max-min) payoff $\frac{1}{2} - c_{x^*}$. This follows directly from the observation that $(x^*, \emptyset) \in Supp(\sigma)$.

4. The equilibrium exhibits a strong correlation between advertising and product quality: the only product that is advertised in equilibrium is the most preferred product.

Our task in this sub-section is to investigate the generality of these properties. The first property is general due to our assumptions that $\phi(x_1, \emptyset) = 0$ for all x_1 and $c_x < c_{x^*}$ for all $x \neq x^*$. This means that if both firms played $(x^*, \emptyset)$, there would be an incentive to deviate to some $(x, \emptyset)$, $x \neq x^*$, because the deviating firm would retain its 50% market share at a lower cost. Thus, it must be the case that firms offer inferior products in equilibrium. The second property is also general. If no marketing devices are employed in equilibrium, then by Lemma 1, $(x, \emptyset) \in Supp(\sigma)$. We have just argued that σ is a mixed-strategy equilibrium so that $Supp(\sigma)$ must contain at least one pair $(x, \emptyset)$, $x \neq x^*$. But these two pure strategies generate different payoffs, $\frac{1}{2} - c_x$ and $\frac{1}{2} - c_{x^*}$, a contradiction.

Our next result demonstrates that the third and fourth properties are general in the sense that there always exists a symmetric equilibrium that satisfies them.

Proposition 2 *There exists a symmetric Nash equilibrium strategy σ such that:*
(i) firms earn a payoff of $\frac{1}{2} - c_{x^*}$ under σ.
(ii) for every $(x, M) \in Supp(\sigma)$, $x = x^*$ or $M = \emptyset$.

Proof. We construct a mixed strategy σ and show that it constitutes a symmetric Nash equilibrium strategy. Let us first construct $Supp(\sigma)$. The first element in $Supp(\sigma)$ is $(x^*, \emptyset)$. Let y^1 be the $\succ$-minimal product in X. By assumption, $(x^*, \emptyset)$ fails to beat $(y^1, \emptyset)$. Add $(y^1, \emptyset)$ to $Supp(\sigma)$. Let m^1 be the least costly marketing device m for which $\phi(y^1, \{m\}) = 1$. By Assumption (*), such a marketing device must exist. Add $(x^*, \{m^1\})$ to $Supp(\sigma)$. This concludes the first step of the construction.

The rest of the construction proceeds iteratively. For some $k \in \{1, ..., |X| - 2\}$, suppose that $Supp(\sigma)$ contains the pairs $(y^1, \emptyset), ..., (y^k, \emptyset)$ and $(x^*, \emptyset), (x^*, \{m^1\}), ..., (x^*, \{m^1, ..., m^k\})$. If $\phi(y, \{m^1, ..., m^k\}) = 1$ for all $y \neq x^*$, then the construction of $Supp(\sigma)$ is complete. Otherwise, let y^{k+1} be the $\succ$-minimal product y for which $\phi(y, \{m^1, ..., m^k\}) = 0$, and add $(y^{k+1}, \emptyset)$ to $Supp(\sigma)$. Let m^{k+1} be the least costly marketing device m for which $\phi(y^{k+1}, \{m\}) = 1$. By Assumption (*), there must exist such a marketing device, and by construction, $m^{k+1} \notin \{m^1, ..., m^k\}$. Add $(x^*, \{m^1, ..., m^k, m^{k+1}\})$ to $Supp(\sigma)$. Assumption (*) guarantees that the iterative process must be terminated after $K \leq |X| - 1$ steps, such that for every $k \leq K$, the strategy $(y^k, \emptyset)$ is beaten by all strategies $(x^*, \{m^1, ..., m^l\})$ with $l \geq k$, and - given our assumption that ϕ is independent of x_2 - by no other strategy in $Supp(\sigma)$.

It remains to assign probabilities to each member of $Supp(\sigma)$. For every $k \in \{1, ..., K\}$ let

$$\sigma(y^k, \emptyset) = 2c_{m^k}$$

and

$$\sum_{l=k}^{K} \sigma(x^*, \{m^1, ..., m^l\}) = 2(c_{x^*} - c_{y^k})$$

In addition, let

$$\sigma(x^*, \emptyset) = 1 - 2[c_{x^*} - c_{y^1} + \sum_{k=1}^{K} c_{m^k}]$$

By our assumptions on costs, all values of $\sigma(\cdot)$ are between zero and one. (Note, however, that if $c_{y^k} = c_{y^{k+1}}$ for some k, then $\sigma(x^*, \{m^1, ..., m^k\}) = 0$, and therefore, strictly speaking, $(x^*, \{m^1, ..., m^k\})$ does not belong to $Supp(\sigma)$.) By construction, the values of $\sigma(\cdot)$ add up to one.

Note that by construction, σ satisfies properties (i) and (ii). First, for every $(x, m) \in Supp(\sigma)$, $x = x^*$ or $M = \emptyset$. Second, since $(x^*, \emptyset) \in Supp(\sigma)$, firms earn a payoff of $\frac{1}{2} - c_{x^*}$ under σ. It thus remains to show that σ constitutes a symmetric Nash equilibrium strategy.

To show this, we first claim that if (y, M) is a best-reply to σ, then so is $(y, \emptyset)$. The expected gain in market share from playing (y, M) instead of $(y, \emptyset)$ is

$$\sum_{m \in M} \sum_{y^k \in X_\phi(m)} \frac{1}{2} \beta_\sigma(y^k)$$

By construction, $\phi(y^k, M) = 1$ if and only if $m^k \in M$, where $m^k \neq m^l$ for $k \neq l$. In addition, $\beta_\sigma(y^k) = 2c_{m^k}$, where m^k is the least costly marketing device m for which $\phi(y^k, \{m\}) = 1$. This means that the expected gain in market share from M cannot be lower than the cost of M. This in turn implies that the expected payoff from $(y, \emptyset)$ cannot be lower than the expected payoff from (y, M).

It follows that in searching for profitable deviations from σ, it suffices to check for strategies of the form $(y, \emptyset)$. By construction, all strategies in $Supp(\sigma)$ generate a payoff of $\frac{1}{2} - c_{x^*}$ against σ. Furthermore, by construction, $x^* \succ y^K \succ \cdots \succ y^1$, and for every y for which $y^{k+1} \succ y \succ y^k$, $c_y \geq c_{y^k}$ and $\phi(y, \{m^k\}) = 1$. This means that $(y, \emptyset)$ generates the same market share as $(y^k, \emptyset)$ and costs no less. Therefore, $(y, \emptyset)$ cannot be a profitable deviation. This concludes the proof. ∎

The result that firms earn max-min payoffs in equilibrium is of interest for several reasons. First, it shows that although consumers' bounded rationality initially creates an opportunity for firms to earn payoffs above the rational-consumer benchmark, competitive forces (which include marketing) eliminate this potential. Second, the equilibrium outcome is Pareto inferior to the rational-consumer benchmark: firms earn the same profits in both cases, while consumers are strictly worse off in the bounded-rationality case. Third, the result turns out to have strong implications for the equilibrium correlation between product quality and marketing, and consequently on consumer conversion rates. We will explore these implications in greater detail in Sub-Section 4.3. At this point, it will suffice to point out that part (ii) of Proposition 2 is not general: there exist equilibria in which firms earn max-min payoffs and yet inferior products are actively marketed.[6]

Symmetric equilibria in which firms earn max-min payoffs have the following interesting property. For every pure strategy (x, M) in the support of the equilibrium strategy σ, each marketing device in M is effective against a distinct set of products that are offered in equilibrium. The equilibrium strategy thus exhibits "marketing efficiency", in the sense that firms employ the minimal marketing devices that are necessary for manipulating consumers' consideration sets. This property was in fact used in the constructive proof of Proposition 2.

Proposition 3 *Let σ be a symmetric Nash equilibrium strategy in which firms earn max-min payoffs. For every $(x, M) \in Supp(\sigma)$ and every $m, m' \in M$, the sets $\{x \in X_\phi(m) : \beta_\sigma(x) > 0\}$ and $\{x \in X_\phi(m') : \beta_\sigma(x) > 0\}$ are disjoint.*

Proof. Assume the contrary - i.e., that there exist $(x, M) \in Supp(\sigma)$ and two marketing devices $m, m' \in M$ such that the two sets $\{x \in X_\phi(m) : \beta_\sigma(x) > 0\}$ and $\{x \in X_\phi(m') : \beta_\sigma(x) > 0\}$ have a non-empty intersection. Then, the marginal contribution of m' to the market share generated by (x, M) is strictly below $\frac{1}{2}\sum_{y \in X_\phi(m')} \beta_\sigma(y)$. Since (x, M) is a best-reply to σ, this implies that $\frac{1}{2}\sum_{y \in X_\phi(m')} \beta_\sigma(y) > c_{m'}$. By the assumption that firms earn max-min payoffs in equilibrium, the strategy $(x^*, \emptyset)$ is a best-reply to σ. It follows that if one of the firms deviates from $(x^*, \emptyset)$ to (x^*, m'), it would earn a payoff in excess of the max-min level, a contradiction. ∎

[6] In the constructive proof of Proposition 2, it is easy to see that when $K > 1$, some of the weight that is assigned to $(x^*, \{m^1\})$, say, can be shifted to a new strategy $(y^2, \{m^1\})$, without upsetting any of the equilibrium conditions.

Suppose that the partitional property described in Proposition 3 holds not only with respect to the products that are offered in equilibrium, but with respect to the grand set of products. That is, assume that the collection $\{X_\phi(m)\}_{m\in\mathcal{M}}$ is a partition of $X\backslash\{x^*\}$. This case fits situations in which there is a pre-existing natural categorization of products (e.g., health versus non-health food products), such that an individual marketing device is effective against a specific category of target products. Under this special case, we are able to provide a complete characterization of the set of symmetric equilibria.

For every $m \in \mathcal{M}$, let $y^*(m)$ denote the $\succ$-minimal product x for which $\phi(x, \{m\}) = 1$. Given a mixed strategy σ, let $\alpha_\sigma(m) = \sum_{M\ni m}\sigma(x, M)$ be the probability that the marketing device m is played under σ.

Proposition 4 *Assume $\{X_\phi(m)\}_{m\in\mathcal{M}}$ is a partition of $X\backslash\{x^*\}$. In any symmetric Nash equilibrium σ:*
(i) firms earn a payoff of $\frac{1}{2} - c_{x^*}$.
(ii) for every $m \in \mathcal{M}$,

$$\begin{aligned} \alpha_\sigma(m) &= 2(c_{x^*} - c_{y^*(m)}) \\ \beta_\sigma(x) &= \begin{cases} 2c_m & if \quad x = y^*(m) \\ 1 - 2\sum_{m\in\mathcal{M}} c_m & if \quad x = x^* \\ 0 & otherwise \end{cases} \end{aligned}$$

Proof. (i) Assume the contrary - i.e., that firms earn more than the max-min payoff $\frac{1}{2} - c_{x^*}$ under some symmetric equilibrium strategy σ. By Lemma 1, $Supp(\sigma)$ contains a strategy of the form (x^*, M). The strategy $(x^*, \emptyset)$ generates the max-min payoff against any strategy. Therefore, it must be the case that $M \neq \emptyset$ for every $(x^*, M) \in Supp(\sigma)$. For every $(x^*, M) \in Supp(\sigma)$ and every $m \in M$, $\frac{1}{2}\sum_{x\in X_\phi(m)}\beta(x) - c_m \geq 0$, with at least one strict inequality for some m^* - otherwise the strategy (x^*, M) could not generate a payoff above the max-min level.

It follows that if $m^* \notin M'$ for some $(x^*, M') \in Supp(\sigma)$, it is profitable to deviate to $(x^*, M' \cup \{m^*\})$. Hence, $m^* \in M$ for all $(x^*, M) \in Supp(\sigma)$. Moreover, because $\frac{1}{2}\sum_{x\in X_\phi(m^*)}\beta(x) - c_{m^*} > 0$, there exists $y_{m^*} \neq x^*$ such that $\beta_\sigma(y_{m^*}) > 0$ and $y_{m^*} \in X_\phi(m^*)$. It must therefore be the case that (x^*, M) beats (y_{m^*}, M') for every $(x^*, M), (y_{m^*}, M') \in Supp(\sigma)$.

Let y^* denote the $\succ$-minimal product among all these products y_{m^*}. If $(y^*, \emptyset)$ is *not* a best-reply to σ, then there must exist m such that $\frac{1}{2}\sum_{x\in X_\phi(m), y^*\succ x}\beta(x) - c_m > 0$ and

$m \in M$ for every $(y^*, M) \in Supp(\sigma)$. But this implies that $m \in M$ for every $(x^*, M) \in Supp(\sigma)$, which means that there is a product y' such that $y^* \succ y'$ and (x^*, M) beats (y', M') for every $(x^*, M), (y', M') \in Supp(\sigma)$, contradicting the definition of y^*. It follows that $(y^*, \emptyset)$ *is* a best-reply to σ.

If a firm deviates from $(y^*, \emptyset)$ to $(x^*, \mathcal{M})$, it will increase its market share by at least $\frac{1}{2}\beta_\sigma(x^*) + \frac{1}{2}(1 - \beta_\sigma(x^*)) = \frac{1}{2}$. Since by assumption $c(x^*, \mathcal{M}) < \frac{1}{2}$, the deviation is profitable, a contradiction. It follows that firms cannot earn more than $\frac{1}{2} - c_{x^*}$ in σ. Since this is the max-min payoff, firms must earn exactly $\frac{1}{2} - c_{x^*}$ in σ.

(ii) First, we claim that for every $m \in M$, $\alpha_\sigma(m) > 0$ and $\beta_\sigma(x) > 0$ for some $x \in X_\phi(m)$. If $\alpha_\sigma(m) = 0$, then when a firm plays $(x, \emptyset)$, where $x \neq x^*$ and $x \in X_\phi(m)$, it earns a payoff $\frac{1}{2} - c_x > \frac{1}{2} - c_{x^*}$, since by assumption $x \notin X_\phi(m')$ for every $m' \neq m$. If $\beta_\sigma(x) = 0$, then it is optimal to set $\alpha_\sigma(m) = 0$, a contradiction.

Second, part (i) implies that for every $m \in M$, $\frac{1}{2}\sum_{x \in X_\phi(m)} \beta(x) \leq c_m$ - otherwise, a firm could play $(x^*, \{m\})$ and earn a payoff above the max-min level. However, if the inequality is strict, firms will find it optimal to set $\alpha_\sigma(m) = 0$. Therefore, $\frac{1}{2}\sum_{x \in X_\phi(m)} \beta(x) = c_m$ for every $m \in \mathcal{M}$. In particular, this means that if $\beta_\sigma(x) > 0$, the strategy $(x, \emptyset)$ must be a best-reply against σ. Denote by $m(x)$ the marketing device which is effective against x. Then, the payoff from $(x, \emptyset)$ is $\frac{1}{2} - c_x - \frac{1}{2}\alpha_\sigma[m(x)]$. Consider a product x satisfying $\beta_\sigma(x) > 0$ and $x \in X_\phi(m)$. If $x \succ y$ for some $y \in X_\phi(m)$, then if a firm deviated to $(y, \emptyset)$ it would earn a payoff of

$$\frac{1}{2} - c_y - \frac{1}{2}\alpha_\sigma[m(x)] > \frac{1}{2} - c_x - \frac{1}{2}\alpha_\sigma[m(x)]$$

in contradiction to the assumption that $\beta_\sigma(x) > 0$. It follows that the only strategy $x \in X_\phi(m)$ for which $\beta_\sigma(x) > 0$ is $y^*(m)$, namely the $\succ$-minimal product in $X_\phi(m)$. And since the payoff from $(y^*(m), \emptyset)$ must be the max-min payoff, it must be the case that $\alpha_\sigma(m) = 2c_{x^*} - 2c_{y^*(m)}$. ■

Thus, when $\{X_\phi(m)\}_{m \in \mathcal{M}}$ is a partition of $X \backslash \{x^*\}$, all symmetric Nash equilibria induce max-min payoffs. Apart from x^*, the only products that are offered in equilibrium are the inferior products in each cell of the induced partition. The more costly the marketing device, the higher the probability with which its inferior target product is offered. The higher the cost of the target product, the lower the probability with which the marketing device is employed. We will have more to say about this example in Sub-Section 4.3.

It is interesting to examine how $(\alpha_\sigma(m))_{m \in \mathcal{M}}$ and $(\beta_\sigma(x))_{x \in X}$ behave with respect to the coarseness of the partition induced by ϕ. For simplicity, let us fix $\sum_{m\in\mathcal{M}} c_m$ and compare two extreme cases: (1) the "Butters" example analyzed in Proposition 1, and (2) the case in which for every $x \neq x^*$ there is a unique marketing device $m(x)$ which is effective against x (i.e., $X_\phi(m) = \{x\}$). The difference between the two cases is that in case 2 marketing devices are specific to a particular target product, while in case 1, marketing is not targeted at all. Since $\sum_{m\in\mathcal{M}} c_m$ is held fixed, $\beta_\sigma(x^*)$ is the same in all symmetric equilibria in both cases. However, in case 2, relative to case 1, some of the weight that the equilibrium strategy assigns to the least preferred product in X is shifted to intermediate quality products. This is a general corollary of Proposition 4: refining the product categorization that is induced by ϕ (while keeping costs fixed in an appropriate sense) results in an upward shift in the equilibrium distribution of product quality.

4.2 Irrelevant alternatives as marketing devices

In the previous sub-section, we assumed that a new product cannot enter the consumer's consideration set without active marketing. The more general case, in which the consideration function also depends on x_2, such that "bare" products can attract serious attention, is more complicated to analyze. In this sub-section we examine a special case in which dominated alternatives serve as marketing devices.

When consumers do not automatically consider all available options, the mere offering of a particular product can have an indirect effect on a firm's market share by drawing attention to the firm and its other products. The items that stores display on their shop front and web retailers put on their homepages can exert a positive externality on other items, by persuading a consumer to enter the store/website and browse its selection. Consumers who enter because of a particular item may end up buying another.

Firms may take this indirect marketing effect into account when designing a product line. Specifically, they may introduce a product even when the demand for it fails to cover its cost. In our market model, since it assumes homogeneous consumer preferences, this effect will take an extreme form. Whenever firms choose to offer multiple products, only one of them (namely, the best product) is ever chosen, whereas the other products are "irrelevant alternatives". Their role is to attract consumers' serious attention to the best product in the product line. In reality, we do not expect such a stark effect, because of consumer heterogeneity. However, our focus on consumer homogeneity allows us to isolate the pure marketing effect of inflating a product line.

Assume that X is a subset of the real numbers - i.e., products can be ordered along a line, as in Hotelling models. Let $\mathcal{M} = X \backslash \{x^*\}$, and let S be the set of all pairs (x, M) such that $x \in X$ and M is a subset of the (strict) $\succ$-lower contour set of x. For every $x \in X$, let $I(x)$ be some neighborhood of x. We define the consideration function as follows: $\phi(x_1, x_2, M_2) = 1$ if $x_1 \in I(y)$ for some $y \in M_2 \cup \{x_2\}$. Our assumption that $c(x, M) < \frac{1}{2}$ for all (x, M) implies that $\sum_{x \in X} c_x < \frac{1}{2}$.

The interpretation is that consideration sets are constructed according to similarity judgments. For each product y, there is a set of products $I(y)$ that resemble it. The consumer is willing to consider substitutes to his default x_1 if the competing firm offers some product which x_1 resembles. Note that the consideration function induces a reflexive binary relation R on X, defined as follows: $x'Rx$ if $x \in I(x')$. This is the similarity relation that underlies the formation of the consideration set.

Unlike the model of the previous sub-section, when a firm offers x *without any marketing devices*, it may still attract consumers' attention, in case the opponent's strategy assigns positive probability to a product $y \in I(x)$. This introduces another layer of complexity into the analysis of equilibrium.

Proposition 5 *Firms earn the max-min payoff* $\frac{1}{2} - c_{x^*}$ *in any symmetric Nash equilibrium.*

The proof is preceded by a pair of lemmas.

Lemma 2 *In any symmetric Nash equilibrium strategy* σ*,* $\beta_\sigma(x) \leq 2c_x$ *for all* $x \neq x^*$*.*

Proof. Assume the contrary. Let x be the $\succ$-minimal product for which $\frac{1}{2}\beta_\sigma(x) > c_x$. Suppose that there exists a pure strategy $(y, M) \in Supp(\sigma)$ such that $y \succ x$ and $x \notin \cup_{m \in M \cup \{y\}} I(m)$. Then, (y, M) fails to beat any pair (x, M'). Suppose that a firm deviates from (y, M) to $(y, M \cup \{x\})$. The deviation does not change the probability of being beaten by some strategy in $Supp(\sigma)$. However, it does change the probability of beating some strategy in $Supp(\sigma)$. Specifically, since $x \in I(x)$, the deviation increases the firm's payoff by at least $\frac{1}{2}\beta_\sigma(x) - c_x > 0$. Therefore, the deviation is profitable. It follows that every $(y, M) \in Supp(\sigma)$ with $y \succ x$ beats every $(x, M') \in Supp(\sigma)$.

Consider some (x, M') in $Supp(\sigma)$. There must be such a strategy, since by assumption, $\frac{1}{2}\beta_\sigma(x) > c_x > 0$. Suppose that a firm deviates from (x, M') to $(x^*, \{x\} \cup M')$. The cost of this deviation is c_{x^*}, whereas the gained market share is at least $\frac{1}{2}\sum_{y \succeq x} \beta(y)$. The reason is that first, $(x^*, \{x\} \cup M')$ beats any pair (x, M'); and second, prior to the deviation, as was shown in the previous paragraph, $(y, M) \in Supp(\sigma)$ with $y \succ x$ had beaten (x, M'), whereas after the deviation no pure strategy can beat$(x^*, \{x\} \cup M')$. In order for this deviation to be unprofitable, we must have

$$\frac{1}{2} \sum_{y \succeq x} \beta(y) \leq c_{x^*}$$

By the definition of x,

$$\frac{1}{2} \beta_\sigma(z) \leq c_z$$

for all $z \prec x$. Adding up these inequalities, we obtain

$$\frac{1}{2} \sum_{y \in X} \beta(y) \leq c_{x^*} + \sum_{z \prec x} c_z < \frac{1}{2}$$

But this implies $\sum_{y \in X} \beta(y) < 1$, a contradiction. ■

Lemma 3 *For every $(x, M) \in Supp(\sigma)$ with $x \neq x^*$, there exists $(x^*, M') \in Supp(\sigma)$ such that (x^*, M') does not beat (x, M).*

Proof. Assume the contrary and let $(x, M) \in Supp(\sigma)$ be a pure strategy that is beaten by all $(x^*, M') \in Supp(\sigma)$. If a firm deviates from (x, M) to some $(x^*, M') \in Supp(\sigma)$, it increases its market share by at least $\frac{1}{2}\beta_\sigma(x^*)$. In order for this deviation to be unprofitable, we must have $\frac{1}{2}\beta_\sigma(x^*) \leq c_{x^*}$. Combined with Lemma 2, we obtain that $\frac{1}{2}\sum_{y \in X} \beta_\sigma(y) \leq \sum_{y \in X} c_y < \frac{1}{2}$, a contradiction. ■

We are now ready to prove Proposition 5.

Proof of Proposition 5. Let us first introduce three pieces of notation. First, define

$$\mathcal{F} \equiv \{M \subseteq \mathcal{M} : (x^*, M) \in Supp(\sigma)\}$$

Second, for every $M \in \mathcal{F}$, define $B(M)$ as the set of products $x \neq x^*$ that satisfy $\beta_\sigma(x) > 0$ and $x \in \cup_{m \in M} I(m)$. Finally, define

$$\Delta(M) \equiv \frac{1}{2} \sum_{z \in B(M)} \beta_\sigma(z) - \sum_{m \in M} c_m$$

Assume that firms earn a payoff above $\frac{1}{2} - c_{x^*}$ under σ. By Lemma 1, $\mathcal{F}$ is nonempty. By Lemma 3, $\beta_\sigma(x) = 0$ for all $x \in I(x^*) \setminus \{x^*\}$. Therefore, in order for a menu (x^*, M), $M \in \mathcal{F}$, to generate a payoff above $\frac{1}{2} - c_{x^*}$, it must be the case that $\Delta(M) > 0$.

Suppose $\mathcal{F} = \{(x^*, M)\}$. Then there must be some $(y, M'') \in Supp(\sigma)$, which is beaten by (x^*, M). But this means that there is no pair (x^*, M') that does not beat (y, M''), contradicting Lemma 3.

Next suppose that $|\mathcal{F}| > 1$. Assume there exists some pair $M, M' \in \mathcal{F}$ such that $B(M) \cap B(M') = \emptyset$. Then, by deviating from (x^*, M') to $(x^*, M' \cup M)$, a firm increases its payoff by $\Delta(M) > 0$, hence the deviation is profitable. Therefore, $B(M) \cap B(M') \neq \emptyset$ for every $M, M' \in \mathcal{F}$. Because $\{I(x)\}_{x \in X}$ is a collection of real intervals, the pairwise intersections imply that $\cap_{M \in \mathcal{F}} B(M) \neq \emptyset$, again contradicting Lemma 3. ■

In the remainder of this sub-section we provide a complete characterization of symmetric equilibria in two special cases.

Identity-based consideration

Our interpretation of consumer behavior in the model of this sub-section is that consumers are willing to consider the product line of a new firm only if it contains a product which is *similar* to the default. An extreme case of this type of behavior is when consumers are reluctant to consider the new firm unless it offers a product which is *identical* to the default - i.e., $I(x) = \{x\}$ for all $x \in X$. For example, think of a consumer who intends to buy a particular digital camera from a certain store. This is his default option. If the consumer happens to pass by a another camera store, he is likely to enquire whether it offers the particular camera he has in mind. The consumer may be more inclined to browse the store's full selection when he is answered, "yes, we have the camera you are interested in, but we also have other cameras which you may like better", than when he is told "no, we don't sell that camera, but we have other cameras, which you may like better".

This extreme case of identity-based similarity turns out to be amenable to complete equilibrium characterization. Let $\alpha_\sigma(x) = \sum_{M \ni x} \sigma(y, M)$ denote the probability that a product x is offered as an *irrelevant* alternative under σ.

Proposition 6 *Suppose that $I(x) = \{x\}$ for all $x \in X$. Then, in any symmetric Nash equilibrium σ, $\beta_\sigma(x) = 2c_x$ and $\alpha_\sigma(x) = 2(c_{x^*} - c_x)$ for all $x \neq x^*$.*

Thus, as an inferior product becomes more costly, it is offered more often as a $\succ$-maximal element in the menu and less often as an irrelevant alternative. We omit the proof because it is very similar to the proof of Proposition 4.

Preference-based consideration

As we remarked in the Introduction, a product may attract the attention of a consumer if it dominates his default along some dimension. When the consumers' true preferences over products reflect a trade-off across several dimensions, some dimensions may be more salient than others. However, this does not necessarily mean that a product that ranks high along one salient dimension is also ranked high when *all* dimensions are taken into account. For instance, the size and price of a car are more easily discerned than its reliability or energy efficiency, and therefore tend to dominate the preference criterion applied in the first stage of the decision procedure, while the other dimensions will come into play only in the second stage. Another example is when consumers have dynamically inconsistent preferences, and they are also naive in the sense that they do not anticipate the change in their preferences. Thus, they choose to enter a store because they expect to make their decision based primarily on one dimension, whereas once inside the store, they give more weight to another dimension.

A real-life example is the addition of "healthier" (but pricier) options, such as salads and fresh fruit, to the McDonald's menu, with the aim of persuading health-conscious customers to enter the restaurant. However, although traffic into McDonald's restaurants has increased, revenues from the unhealthy "Dollar Menu" have significantly increased at the expense of the healthier options (see Warner (2006)). This example demonstrates that when a consumer is initially reluctant to consider a firm's main product, certain additional products can provide a justification for considering the main product, thus creating the possibility that the consumer will eventually choose it despite his initial reluctance.

To capture these situations, assume that the elements in X are ordered $x_1 < x_2 < \cdots < x_n$ such that $I(x_1) \subseteq \cdots \subseteq I(x_n)$. Recall that ϕ induces the binary relation R on X, yRx if $x \in I(y)$. Our assumption on $I(\cdot)$ implies that R is complete and transitive. Hence, we may interpret R in this case as a ranking of products according to some salient dimension. The following result characterizes symmetric equilibria in this case.

Proposition 7 *Let σ be a symmetric Nash equilibrium. For every $(x, M) \in Supp(\sigma)$, $|M| \leq 1$, with strict inequality whenever $x \neq x^*$.*

Proof. Suppose that $Supp(\sigma)$ contains a pure strategy (x, M) with $|M| > 1$. By our assumption on X and $\{I(x)\}_{x \in X}$, there exists $y \in \{x\} \cup M$ such that $I(z) \subseteq I(y)$ for any $z \in \{y\} \cup M$. This means that the strategy $(x, \{y\})$ attains the same market share as (x, M), yet it is less costly, a contradiction. It follows that $|M| \leq 1$ for every $(x, M) \in Supp(\sigma)$.

Suppose next that there exists a strategy $(x, \{y\}) \in Supp(\sigma)$, where $x \neq x^*$. This means that $I(x) \subset I(y)$ - otherwise it would be profitable to deviate from $(x, \{y\})$ to $(x, \emptyset)$. Therefore, $x^* \in I(y)$. Now consider the strategy $(x^*, \{y\})$. Note that $x^* \neq y$ because $y \prec x$. This strategy beats any strategy that $(x, \{y\})$ beats, but it also beats $(x, \{y\})$ because $x \in I(y)$ and $x^* \succ x$. Therefore, given that the strategy $(x, \{y\})$ is weakly more profitable than the strategy $(x, \emptyset)$ against σ, the strategy $(x^*, \{y\})$ must be strictly more profitable than the strategy $(x^*, \emptyset)$ against σ. Since the latter strategy always generates at least the max-min payoff against any strategy, it follows that $(x^*, \{y\})$ generates a payoff strictly above the max-min level, in contradiction to Proposition 5. ■

Proposition 7 says that when the consideration function is based on a preference relation, firms use the marketing device of a single irrelevant alternative in symmetric Nash equilibrium only to promote the most preferred product x^*.

Distinguishing between products and marketing devices

Our ability to distinguish between products and marketing devices in this sub-section relied on the assumption of homogeneous preferences. We would not be able to make this distinction in an environment with heterogeneous preferences. However, there is a natural way of reformulating the model of this sub-section in a way that avoids making that distinction. According to this reformulation, a pure strategy for a firm is a *menu*, namely a non-empty subset of products $A \subseteq X$. The cost of a menu A is $\sum_{x \in A} c_x$. Let $b(A)$ denote a consumer's most preferred product in A. Let A_1 denote the menu associated with the firm to which the consumer is initially assigned, and let A_2 denote the menu offered by the rival firm. The consumer's procedure then implies that he chooses $b(A_2)$ from A_2 if $b(A_2) \succ b(A_1)$ and $b(A_1) \in I(y)$ for some $y \in A_2$, and he chooses $b(A_1)$ from A_1 otherwise.

4.3 Consumer conversion

In several instances of the market model - namely, the "Butters" advertising example in Sub-Section 4.1 and the preference-based consideration model of Sub-Section 4.2 - we saw that in a symmetric equilibrium, firms use active marketing only to promote the most preferred product x^*. As we already mentioned, this is not a general property. Instead, there is a weaker property that captures the correlation between product quality and marketing in symmetric equilibria that induce max-min payoffs.

Definition 1 (Effective Marketing Property) *A mixed strategy σ satisfies the effective marketing property if for every $(x, M), (x', M') \in Supp(\sigma)$, if $\phi(x, x', M') = 1$ and $\phi(x, x', \emptyset) = 0$, then $x' \succ x$ or $x = x^*$.*

The effective marketing property means that whenever a consumer considers a new product thanks to the marketing strategy that accompanies that product, he ends up buying it, unless his default product is x^*.

Proposition 8 *Let σ be a symmetric Nash equilibrium strategy that induces max-min payoffs. Then, σ satisfies the effective marketing property.*

Proof. Let $(x, M), (x', M') \in Supp(\sigma)$, $\phi(x, x', M') = 1$, $\phi(x, x', \emptyset) = 0$, $x \neq x^*$ and yet $x \succeq x'$. For every strategy (x', M'), let $B(x', M')$ denote the set of strategies in $Supp(\sigma)$ that (x', M') beats. Recall that the set of strategies that beat (x', M') is independent of M'. In order for (x', M') to be a best-reply to σ, it must be weakly preferred to $(x', \emptyset)$, and therefore satisfy the following inequality:

$$\sum_{(y,N) \in B(x',M') \setminus B(x',\emptyset)} \sigma(y, N) \geq 2 \sum_{m \in M'} c_m$$

By the assumption that firms earn max-min payoffs under σ, the strategy $(x^*, \emptyset)$ is a best-reply to σ, and moreover, $B(x^*, \emptyset) = \emptyset$. Suppose that a firm deviates to (x^*, M'). In order for this deviation to be unprofitable, the following inequality must hold:

$$\sum_{(y,N) \in B(x^*,M')} \sigma(y, N) \leq 2 \sum_{m \in M'} c_m$$

Because $x^* \succ x'$, $B(x', M') \setminus B(x', \emptyset) \subseteq B(x^*, M')$. Moreover, since $\phi(x, x', M') = 1$ and $x^* \succ x \succeq x'$, the inclusion is strict. Therefore,

$$\sum_{(y,N) \in B(x^*,M')} \sigma(y, N) > \sum_{(y,N) \in B(x',M')} \sigma(y, N)$$

which contradicts the combination of the preceding pair of inequalities. ■

The effective marketing property is a result that characterizes *consumer conversion rates* - that is, the probability that a consumer will switch to the new product conditional on having considered it. Of course, the result that the conversion rate is essentially 100% is unrealistic, and clearly relies on the equally unrealistic assumption of consumer homogeneity. We view the result as a useful theoretical benchmark for richer, more pertinent theories of conversion rates, that incorporate consumer heterogeneity, among other things.

For some specifications of the model, we can also characterize the unconditional probability that consumers switch a supplier. Consider, for instance, the case studied in Sub-Section 4.1 where the collection $\{X_\phi(m)\}_{m\in\mathcal{M}}$ is a partition of $X\backslash\{x^*\}$. Proposition 4 characterized $(\alpha_\sigma(m))_{m\in\mathcal{M}}$ and $(\beta_\sigma(x))_{x\in X}$. By the effective marketing property, the probability that a consumer whose default is $x \neq x^*$ will switch a supplier is $\alpha_\sigma(m(x))$, where $m(x)$ denotes the marketing device that satisfies $\phi(x, \{m\}) = 1$. Therefore, by Proposition 4, the overall switching rate is

$$\sum_{x\neq x^*} \beta_\sigma(x)\alpha_\sigma(m(x)) = \sum_{m\in\mathcal{M}} \beta_\sigma(y^*(m))\alpha_\sigma(m) = 4\sum_{m\in\mathcal{M}} c_m \cdot (c_{x^*} - c_{y^*(m)})$$

where $y^*(m)$ denotes the least preferred product against which m is effective.

Thus, the switching rate increases with the cost of marketing and decreases with the cost of inferior products. Note that the switching rate is equal to the expected cost of marketing devices under σ. This follows from the observation that the probability that a marketing device is employed by a given firm is $\alpha_\sigma(m)$, and the cost of the marketing device is $c_m = \frac{1}{2}\beta_\sigma(y^*(m))$.

In the case of identity-based consideration analyzed in Sub-Section 4.2, the effective marketing property implies that the overall switching rate is

$$\sum_{x\neq x^*} \beta_\sigma(x)\alpha_\sigma(x) = 4\sum_{x\neq x^*} c_x \cdot (c_{x^*} - c_x)$$

Thus, in this case the switching rate behaves non-monotonically in the cost of inferior products. This is due to the double role of products as both consumption alternatives and marketing devices.

4.4 Can firms attain collusive profits in equilibrium?

Recall the model of Sub-Section 4.1, and imagine a scale that measures consumers' resistance to considering new alternatives. At one end of the scale we have the fully rational consideration function which always yields the feasible set. Suppose that at the other end of the scale we place the consideration functions for which $\{X_\phi(m)\}_{m\in\mathcal{M}}$ constitutes a partition of $X \setminus \{x^*\}$. At both ends of this scale, the fully rational one and the boundedly rational one, we saw that firms necessarily earn the max-min payoff in symmetric Nash equilibrium. Intuitively, one would expect the competition between firms to be fiercer, the closer we move to the rational end of the scale. According to this intuition, firms would not be able to make collusive profits when the consideration set becomes more likely to coincide with the objectively feasible set.

This intuition turns out to be false, as the following counter-example demonstrates. Let $X = \{111, 100, 010, 001\}$ and $x^* = 111$. Let $\mathcal{M} = \{1, 2, 3\}$ and assume the following consideration function ϕ: $X_\phi(1) = \{001\}$, $X_\phi(2) = \{100\}$ and $X_\phi(3) = \{010\}$. Thus, $\{X_\phi(m)\}_{m\in\mathcal{M}}$ constitutes a partition of $X \backslash \{x^*\}$. Let $c_{111} = \frac{1}{3}$, and let $c_m = c_x = \bar{c} < \frac{1}{30}$ for all $m \in \mathcal{M}$ and $x \neq x^*$. By Proposition 4, in every symmetric Nash equilibrium, firms earn the max-min payoff.

Now consider modifying the consumers' consideration function into ϕ', such that $X_\phi(1) = \{010, 001\}$, $X_\phi(2) = \{100, 001\}$ and $X_\phi(3) = \{100, 010\}$. This modification has a natural interpretation. Each product may have up to three attributes. The most preferred product has all three attributes. A marketing device m is interpreted as an ad that focuses on the m-th attribute. If the consumer's default product lacks that attribute, the ad persuades him that he should consider the new product. Note that $\{X_{\phi'}(m)\}_{m\in\mathcal{M}}$ is a cover, but not a partition of $X \setminus \{x^*\}$.

It can be shown that the modified consideration function generates a continuum of symmetric equilibria, in which the support of the equilibrium strategy consists of the strategies $(111, \{1\})$, $(111, \{2\})$, $(111, \{3\})$, $(100, \emptyset)$, $(010, \emptyset)$ and $(001, \emptyset)\}$, and firms earn payoffs *above* the max-min level. This example is a counter-part to Proposition 2: it demonstrates that the model of Sub-Section 4.1 may have symmetric Nash equilibria in which firms attain collusive profits.[7]

Although the consideration function that gives rise to the counter-example is natural, the restriction on the cost function is non-generic. Is it true that for generic cost functions, any symmetric Nash equilibrium induces max-min payoffs? This is an open problem we leave for future research. At any rate, our final result in this section demonstrates that when costs are sufficiently low, equilibrium payoffs are equal to the max-min level.

Proposition 9 *If $c(x^*, \mathcal{M}) < 1/(2^{|\mathcal{M}|+1} + 2)$, then firms earn the max-min payoff in any symmetric Nash equilibrium.*

[7] A similar example can be constructed for an extended version of the model of Sub-Section 4.2, in which $\{I(x)\}_{x\in X}$ is *not* a collection of real intervals. Let X be the set of all 3-digit binary numbers except 000, and assume that $I(x)$ is the set of products that share at least two digits with x. There exist cost parameters that give rise to a continuum of symmetric equilibria in which firms earn payoffs above the max-min level.

Proof. Assume, by contradiction, that $c(x^*, \mathcal{M})) < 1/(2^{|\mathcal{M}|+1} + 2)$, and yet firms earn payoffs above the max-min level in some symmetric Nash equilibrium σ.

We first claim that

$$\frac{1}{2} \sum_{x \prec x^*} \beta_\sigma(x) < c(x^*, \mathcal{M}) \tag{2}$$

To see why this is true, consider some $(x^*, M) \in Supp(\sigma)$. By Lemma 1, $Supp(\sigma)$ must contain such a strategy. By our assumption that firms earn an expected payoff above the max-min, (x^*, M) must beat some other strategy $(x, M') \in Supp(\sigma)$. Define $B_v(M) \equiv \{x \prec x^* : \phi(x, x^*, M) = v\}$. Note that $B_0(M) \cup B_1(M) = \{x \in X \mid x \prec x^*\}$. For any $(x^*, M) \in Supp(\sigma)$, it must be the case that

$$\frac{1}{2} \sum_{x \in B_0(M)} \beta_\sigma(x) \leq c(x^*, \mathcal{M}) - c(x^*, M)$$

Otherwise, it is profitable to deviate from (x^*, M) to $c(x^*, \mathcal{M})$. In addition, it must be the case that

$$\frac{1}{2}\sigma(x^*, M) + \frac{1}{2} \sum_{x \in B_1(M)} \beta_\sigma(x) \leq c(x^*, M) - c(x', M')$$

for some $(x', M') \in Supp(\sigma)$ that is beaten by (x^*, M). Otherwise, it is profitable to deviate from (x', M') to (x^*, M). Summing over the last two inequalities, we obtain inequality (2).

Since every $(x^*, M) \in Supp(\sigma)$ must beat some $(x', M') \in Supp(\sigma)$, it must be the case that

$$\frac{1}{2}\sigma(x^*, M) \leq c(x^*, M) - c(x', M') < c(x^*, \mathcal{M})$$

Otherwise, it would be profitable to deviate from (x', M') to $(x^*, \mathcal{M})$. The number of strategies of the form (x^*, M) in $Supp(\sigma)$ is at most $2^{|\mathcal{M}|}$. Summing over all these strategies, we obtain

$$\frac{1}{2} \sum_{M} \sigma(x^*, M) = \frac{1}{2}\beta_\sigma(x^*) < 2^{|\mathcal{M}|} \cdot c(x^*, \mathcal{M})$$

Combined with the inequality (2), we obtain

$$1 < (2^{|\mathcal{M}|+1} + 2) \cdot c(x^*, \mathcal{M})$$

a contradiction. ■

Proposition 9 can actually be proven for a more general set-up than the one described in Section 3. This result continues to hold when we allow ϕ to be a function of M_1, when ϕ is not "separable", and when the cost of a strategy (x, M) is not additive. There are only two crucial assumptions for the result: (i) the existence of a marketing strategy M^* such that $\phi(x, M, x^*, M^*) = 1$ for all x, M, and (ii) the cost of strategies is bounded from above by $\frac{1}{2}$.

Still, Proposition 9 is somewhat unsatisfactory for the following reason. When costs are small, the probability that x^* is offered is close to one, as can easily be seen from inequality (2). Thus, a max–min payoff result that holds only when costs are very small takes some of the sting out of the distinction between the coincidence of the market outcome with the rational-consumer benchmark and the coincidence of industry profits with the rational-consumer benchmark.

5 Discussion

In this section we discuss the relation between the consideration sets model and various branches of related literature.

5.1 Choice-theoretic aspects of the model

In this sub-section we examine some revealed-preference properties of the consideration-sets model, and compare it to related models in the choice theoretic literature.

Let $\succ^*$ be a binary relation over $X \times 2^{\mathcal{M}}$ defined as follows: $(y, N) \succ^* (x, M)$ if $\phi(x, M, y, N) = 1$ and $y \succ x$. This is the (strict) revealed preference relation induced by the $(\succ, \phi)$ procedure. This binary relation may violate *transitivity*. To see why, consider the following example. Assume $x'' \succ x' \succ x$, $\phi(x, M, x', M') = 1$, $\phi(x', M', x'', M'') = 1$ and $\phi(x, M, x'', M'') = 0$. Then, $(x', M') \succ^* (x, M)$ and $(x'', M'') \succ^* (x', M')$, yet $(x'', M'') \not\succ^* (x, M)$. Likewise, it can be shown that the *weak* revealed preference relation induced by the choice procedure may be incomplete as well as intransitive.

The revealed preference relation does satisfy certain rationality properties. First, although $\succ^*$ may violate transitivity, it does not contain cycles of any length. In addition, $\succ^*$ satisfies the following property: $(y, N) \succ^* (x, M)$ implies $(x, M') \not\succ^* (y, N')$ for all $M', N' \subseteq \mathcal{M}$. That is, marketing cannot reverse the consumer's revealed preferences over products. In particular, when the two extended alternatives are simply the same product in two different guises, the consumer never strictly prefers one extended alternative to another. The reason is that in our model, marketing can manipulate consumers' perception of the feasible set, but it does not manipulate their preferences.

Which properties of $\succ^*$ characterize the consideration sets procedure? That is, can we state axioms on $\succ^*$ that will be satisfied if and only if there exist a linear ordering $\succ$ on X and a consideration function $\phi : S^2 \longrightarrow \{0,1\}^2$ such that $(y, N) \succ^* (x, M)$ if and only if $\phi(x, M, y, N) = 1$ and $y \succ x$? Without imposing any restriction on the consideration function ϕ, this question is trivial, since we can attribute the entire choice behavior to ϕ by simply letting $\phi(x, M, y, N) = 1$ if and only if $(y, N) \succ^* (x, M)$. Choice-theoretic characterization of the consideration sets procedure under some restrictions on ϕ (such as those imposed in Sections 3 and 4) would take us beyond the scope of the present paper.

Masatlioglu and Nakajima (2008) independently conduct a choice-theoretic analysis of a more general choice procedure than ours, which they call "Choice by Iterative Search" (CIS). A consumer who follows this procedure begins with some exogenously given default option r, taken from the feasible set B. Given this default, the consumer constructs a consideration set $\Omega(B, r) \subseteq B$. The consumer chooses the best alternative in $\Omega(B, r)$ according to a complete preference relation $\succ^*$ defined on the grand set of alternatives X. If $\max_{\succ} \Omega(B, r) = r$, the procedure is terminated and the consumer chooses r. If $\max_{\succ} \Omega(B, r) \neq r$, then the consumer constructs another consideration set $\Omega[B, \max_{\succ} \Omega(B, r)]$ and picks his most preferred alternative from this set. The procedure is iterated until the consumer picks some alternative y that satisfies $\max_{\succ} \Omega(B, y) = y$. The CIS procedure is characterized by the mapping Ω and the preference relation $\succ$.

Our choice procedure is a special case of the CIS model. Given a pair of extended alternatives $(x_1, M_1), (x_2, M_2)$, let $B = \{(x_1, M_1), (x_2, M_2)\}$, $r = (x_1, M_1)$, $\Omega(B, r) = B$ if $\phi(x_1, M_1, x_2, M_2) = 1$, and $\Omega(B, r) = \{r\}$ if $\phi(x_1, M_1, x_2, M_2) = 0$. The preference relation $\succ^*$ coincides with $\succ$ in our model. Masatlioglu and Nakajima (2008) show that the CIS model induces an extended choice function (a mapping from pairs, consisting of a set B and a default r, to an element in B) which is fully characterized by two properties, which they call "Anchor Bias" and "Dominating Anchor Bias".[8]

The notion of consideration sets is also related to the idea of "shortlisting". A decision maker who faces a large choice set may simplify his decision problem by first eliminating a subset of alternatives that are dominated according to some incomplete

[8]Masatlioglu and Nakajima also provide a choice-theoretic characterization when the default is not observed, but has to be inferred from observations. A choice correspondence satisfies a property called "Bliss-Point" if, and only if, there exist a preference relation over alternatives $\succ^*$ and a consideration set mapping Ω, such that for every $B \subseteq X$, each element chosen from B is selected by the CIS procedure $(\Omega, \succ^*)$ for some default.

preference relation, and then applying a complete preference relation to the remaining set. Manzini and Mariotti (2007) provide a choice-theoretic characterization of this procedure.

The intuitive difference between the two models is that although they both apply a pair of binary relations in sequence, the shortlisting model uses the first stage to *shrink* the choice set, whereas the consideration-sets model uses the first stage to *expand* it. Note that unlike the consideration-sets procedure, the shortlisting model does not involve an explicit default alternative. For this reason, a straightforward comparison between the two models is impossible. Nevertheless, certain partial comparisons are feasible.

First, we can compare the choice behavior induced by the two models when the default is held fixed. Consider the specification of our model studied in Sub-Section 4.2, in which marketing devices are inferior alternative in a firm's product line. Let $X = \{x, y, z\}$ such that $y \succ z \succ x$, $\phi(z, x, \emptyset) = 1$, $\phi(z, y, \{x\}) = 1$ and $\phi(z, y, \emptyset) = 0$. When the consumer faces the extended choice problem $\{(z, \emptyset), (x, \emptyset)\}$, he considers both z and x, yet chooses z because he finds it better than x. When the extended choice problem is $\{(z, \emptyset), (y, \emptyset))$, the consumer chooses z despite its inferiority to y, because he does not consider y. Finally, when the extended choice problem is $\{(z, \emptyset), (y, \{x\})\}$, the consumer chooses y because the irrelevant alternative x serves as a "door opener" that convinces him to consider y in addition to the default z. This choice behavior violates a necessary condition for shortlisting called "expansion": if an element is chosen from two sets, then it should be chosen from their union as well (see Manzini and Mariotti (2007), p.1828).

Second, consider the special case in which the binary relations that are employed in both stages of the shortlisting model are complete and transitive. Then, the shortlisting model is reduced to standard rational choice. Compare this with the case in our model where the consideration function induces a complete and transitive binary relation P on extended alternatives, defined as follows: $(y, N)P(x, M)$ if $\phi(x, M, y, N) = 1$. In this case, the consumer chooses his default x_1, unless $(x_2, M_2)P(x_1, M_1)$ *and* $x_2 \succ x_1$, in which case he chooses x_2. Thus, in order for the consumer to switch from the default (x_1, M_1) to the new alternative (x_2, M_2), the latter must be ranked above the former according to *two* preference relations. As Masatlioglu and Ok (2005) showed, this sort of behavior is consistent with choosing according to an *incomplete* preference relation over $X \times 2^{\mathcal{M}}$, where a new alternative is chosen over the default only if it is strictly better according to this incomplete preference relation.

An crucial difference between the consideration-sets procedure and both the CIS and shortlisting models is that our model imposes more structure on the set of outcomes, in the form of the distinction between products and marketing strategies. Rubinstein and Salant (2008) study a choice model that involves a related distinction between "alternatives" and "frames". In their model, the frame accompanies the entire choice set rather than an individual alternative. Of course, one can translate our concept of a frame into theirs by taking the profile of marketing strategies to be the frame that accompanies the choice set. Rubinstein and Salant provide necessary and sufficient conditions for rationalizing a choice function (defined over framed choice problems) with a (possibly incomplete) preference relation defined over the set of alternatives.

5.2 Persuasive, complementary and informative advertising

Models of advertising in economics typically make one of the following assumptions (see Bagwell (2007)): (*i*) advertising changes the utility function from consumption (advertising is "persuasive"); (*ii*) advertising enters into the utility function as an argument (advertising is "complementary" to consumption); and (*iii*) advertising does not affect the utility function but it affects the consumer's beliefs (advertising is "informative"). In this sub-section we try to relate our model to this categorization.

Persuasive and complementary advertising

Recall from Sub-Section 5.1 that if a consumer in our model switches from a default product x to a competing product y as a result of the marketing of y, then no set of marketing devices would cause the consumer to switch from y to x. Hence, our framework cannot accommodate any model of persuasive or complementary advertising that allows such preference reversals. This raises the following question: can consumer behavior in our framework always be modelled as some form of persuasive or complementary advertising?

Persuasive advertising means that the consumer's preference relation over products is indexed by (M_1, M_2), the profile of marketing strategies. That is, a consumer is characterized by a profile of *weak* preference relations $(\succsim_{(M_1,M_2)})_{(M_1,M_2)\in 2^{\mathcal{M}}}$ over X. The question is, can we associate with any pair $(\succ, \phi)$ a profile of preference relations $(\succsim_{(M_1,M_2)})_{(M_1,M_2)\in 2^{\mathcal{M}}}$ such that for every tuple (x_1, M_1, x_2, M_2) the consumer's choice is $\max_{\succsim_{(M_1,M_2)}}\{x_1, x_2\}$?

The answer is negative, due to the possible "intransitivity" of ϕ. In particular, it is possible that $\phi(x, M_1, y, M_2) = 1$, $\phi(y, M_1, z, M_2) = 1$ and yet $\phi(x, M_1, z, M_2) = 0$. If $z \succ y \succ x$, the consumer will choose y over x and z over y, but he will not choose z over x. Therefore, since we have held (M_1, M_2) constant, no preference relation $\succsim_{(M_1,M_2)}$

can rationalize this choice behavior.

Chioveanu (in press) analyzes an extension of Varian's model of sales (Varian (1980)), in which some consumers rationally perform price comparisons at no cost, while other consumers are loyal to firms they are initially assigned to, where loyalty means that they do not perform any price comparison. Chioveanu assumes that the fraction of consumers who are loyal to a given firm in this sense is a function of the profile of advertising expenditures in the industry. Although Chioveanu refers to this advertising technology as "persuasive", it does not fall into the definition of persuasive advertising given above. Instead, the way Varian and Chioveanu model customer loyalty and persuasive advertising fits in well with our model: a consumer is "loyal" to a firm if his consideration set consists of the firm's product only.

Advertising is complementary if the revealed choices of the consumer can be rationalized by a single preference relation over the extended set of alternatives $X \times 2^{\mathcal{M}}$. As we remarked in Sub-Section 5.1, our model can induce choice behavior that cannot be rationalized by standard preferences over $X \times 2^{\mathcal{M}}$. Hence, our model accommodates choice behavior that cannot be captured by a model of complementary advertising.

Informative advertising

Informative advertising typically takes two forms. First, in a search-theoretic environment, advertising can reduce the search costs that the consumer needs to incur in order to add a product to his choice set (in extreme cases, such as in Butters (1977), costs fall from being infinitely high to being zero). Second, advertising can cause the consumer to update his beliefs about the quality of the product, either because the advertising message contains verifiable data or because it acts as a Spencian signalling device.

The behavioral comparison between our model and informative advertising is subtle, because the latter approach assumes that the consumer has rational expectations about the distribution of alternatives he is facing, a component that is absent from our model. However, any model of informative advertising would necessarily display the following *monotonicity* property. If an advertising message convinces the consumer to consider a new product when his default is x, it should also convince him to consider the new product when his default is inferior to x according to his preferences. The models analyzed in Sub-Sections 4.1 and 4.2 typically violate this type of monotonicity.

Conclusion

Thus, the consideration-sets model departs from the trinity of persuasive, comparative and informative advertising. In our model, the role of marketing is to "persuade to consider", and this role is related to, but distinct from these three conventional theories. Finally, recall that our model incorporates other marketing activities than advertising, including packaging, determination of payoff-irrelevant product characteristics, search engine optimization and design of product lines.

5.3 Related marketing literature

The marketing literature has long recognized that the consumption decision follows a two-step decision process (for extensive surveys of this literature, see Alba, Hutchinson and Lynch (1991) and Roberts and Lattin (1997)). Consumers first form a small set of options that they will consider for their consumption decision. They then evaluate the options in this set and choose the one they prefer the most. Whether or not an alternative is included in the consideration set may depend on factors other than the consumer's preferences.

Empirical evidence for this two-stage procedure is not trivial to gather, because the first stage is hard to observe. In a study of laundry detergent purchases, Hoyer (1984) reports that the median number of packages that consumers closely examined, as they browsed the relevant supermarket shelf, was one. Thus, even if new, superior brands were displayed on the shelf, it is unlikely that they would have been considered by the consumer, unless they were promoted.

Shum (2004) presents evidence that is consistent with the view that marketing attempts to weaken consumers' reluctance to consider new products. He carries out counterfactual experiments which demonstrate that uninformative advertising may be at least as effective as price discounts in stimulating a purchase of a new brand.

Alba et al. (1991) emphasize the important role that memory plays in the formation of consideration sets. First, many purchasing decisions are made without having the feasible alternatives physically present (e.g., deciding on a restaurant for dinner). Second, even when the available options are displayed to the consumer, the display is often complex (e.g., financial products, sophisticated electrical appliances) or overwhelmingly varied (e.g., breakfast cereals or salad dressing in a supermarket). In these circumstances, consumers rely on memory to a large extent. This implies that a preferred option may be ignored if it is not easily retrieved from memory.

For example, Nedungadi (1989) studied the effect of uninformative advertising on choice of fast food restaurant. Subjects were told that they would be given a coupon for a fast food restaurant of their choice. On the premise that the experimenter had only a limited variety of coupons available to him, subjects were asked to name their most preferred restaurant and list all other restaurants for which they would accept a coupon. In one treatment, before subjects provided the names, they were exposed to an ad that mentioned a local sandwich shop (without any information on this shop's menu). Subjects in the control treatment were not exposed to this ad. Nedungadi found that while most subjects in the control treatment listed mainly hamburger restaurants, a significant proportion of subjects in the advertising treatment named a well-known sandwich chain - *different than the one which was advertised* - as their most preferred choice. Thus, even though some subjects preferred sandwiches to hamburgers, the former was unlikely to be chosen simply because it was not easily recalled when the task was to choose a fast food restaurant.

Memory also plays a role in the choice between an existing brand of an incumbent firm and a new competing brand of an entrant. The likelihood of choosing the new product depends on the ease with which this product will be retrieved whenever the consumer considers making a purchase from the product class to which it belongs. Zhang and Markman (1998) propose that the likelihood of remembering a new brand is influenced by the way its attributes compare with those of the incumbent brand. Specifically, the authors provide experimental evidence suggesting that consumers are more likely to recall a new brand if its advertised attributes are comparable with the attributes of the incumbent brand along a common dimension (i.e. the differences between the two brands are *alignable*). Moreover, the authors demonstrate that a superior new brand may not be chosen if its good attributes are hard to align with those of the incumbent brand. In a similar vein, a recent study by Chakravarti and Janiszewski (2003) presents experimental evidence suggesting that when people are asked to select an alternative from a large set of heterogeneous alternatives, they tend to simplify their decision problem by focusing on a small subset of "easy-to-compare" options having alignable attributes.

6 Concluding remarks

This paper introduces the concept of consideration sets into economic modeling and develops its implications in the context of a competitive market model. As such, it contributes to a growing theoretical literature on market interactions between profit-maximizing firms and boundedly rational consumers. Rubinstein (1993) analyzes monopolistic behavior when consumers differ in their ability to understand complex pricing schedules. Piccione and Rubinstein (2003) study intertemporal pricing when consumers have diverse ability to perceive temporal patterns. Spiegler (2006a,b) analyzes markets in which profit-maximizing firms compete over consumers who rely on naive sampling to evaluate each firm. Shapiro (2006) studies a model in which firms use advertising to manipulate the beliefs of consumers with bounded memory. DellaVigna and Malmendier (2004), Eliaz and Spiegler (2005, in press), and Gabaix and Laibson (2006) study interaction with consumers having limited ability to predict their future tastes. Mullainathan, Schwartzstein and Shleifer (2008) study the role of uninformative advertising when consumers apply "coarse reasoning". For a field experiment that quantifies the effects of various marketing devices in terms of their price-reduction equivalent, See Bertrand, Karlan, Mullainathan, Shafir and Zinman (2008).

We hope to extend our market model in various directions. An important challenge is to incorporate price setting into the model. This extension is not mechanical. On one hand, it makes sense to assume that the consideration function depends on prices. On the other hand, there is no single obvious way to do this. One way to proceed is to assume that an extended alternative is a product x and a price p, such that prices play the role of marketing devices (only here we assume that a firm chooses a single marketing device). Clearly, consumer preferences satisfy $(z, p') \succ (z, p)$ whenever $p' < p$. A natural specification of ϕ in this case is $\phi(x, p, y, p') = 1$ if and only if $p' < p$. Alternatively, we can assume that $\phi(x, p, y, p') = 1$ if and only if $x = y$ and $p' < p$. One challenge is to characterize the structure of equilibrium price dispersion and product variety in such an extended model.

Another important extension of the model is in the direction of consumer heterogeneity. Since consumers in our model are characterized by two primitives, $\succsim$ and ϕ, heterogeneity may exist in both dimensions. Here we make do with a brief comment on a somewhat surprising effect of heterogeneity in ϕ. Formally, heterogeneity means that ϕ is a function that gets values in $[0, 1]$, rather than in $\{0, 1\}$. Let us revisit the example of Sub-Section 4.1 in which ϕ induces a partition on the set of products. Suppose that we mix the population of consumers with a small group of rational consumers. The max-min payoff is unaltered by this modification. However, firms earn a payoff above the max-min level in equilibrium. The reason is that if there are not too many rational consumers, $\beta_\sigma(x^*)$ will be strictly between zero and one. This means that inferior products will be offered with positive probability. But thanks to the presence of rational consumers, the strategy $(x^*, \emptyset)$ generates a market share above 50%, and therefore a payoff above the max-min level. This is another instance of the anomalous effect pointed out in Sub-Section 4.4: making the population of consumers "more rational" does not imply that industry profits will necessary go down.

Finally, it would be interesting to study market models in which the consideration set is also a function of M_1, the set of marketing devices that accompany the default. This would allow us to examine situations where a product's marketing can dissuade consumers from considering competing alternatives, as in example 2.4, say. To get a glimpse into such an extension, consider a generalization of this example. Suppose there exists a compete and transitive binary relation P over $X \times M$, such that $\phi(x, M, y, N) = 1$ if and only if $(y, N)P(x, M)$.[9] This family of consideration functions also captures situations in which the consumer's decision whether to consider a new product depends on a superficial preference relation, which may depend on how each product is presented (its packaging or image) in addition to its actual content. The superficial preference criterion P and the consumer's true preference relation $\succ$ may or may not overlap. It can be shown that under this class of consideration functions, firms earn the max-min payoff in any symmetric Nash equilibrium. However, unlike the environments studied in Section 4, this result does not give rise to an effective marketing property, because of the preventive aspect of marketing captured by the assumption that the consideration function depends on the marketing devices that accompany the default.

References

[1] Alba, J W., J.W. Hutchinson, and J.G. Lynch, Jr., "Memory and Decision Making," in H. Kassarjian and T. Robertson (Eds.) *Handbook of Consumer Behavior*. Englewood Cliffs, NJ: Prentice Hall, 1991, 1-49.

[2] Butters, G. (1977): "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies* **44**, 465-491.

[3] Bertrand, M., S. Mullainathan, E. Shafir and J. Zinman (2008): "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment," Mimeo.

[4] Chakravarti, A. and C. Janiszewski (2003): "The Influence of Macro-Level Motives on Consideration Set Composition in Novel Purchase Situations," *Journal of Consumer Research* **30**, 244-258.

[5] Chioveanu, I. (in press): "Advertising Brand Loyalty and Pricing," *Games and Economic Behavior*.

[9] in example 2.4, the relation P is defined as follows: $(y, N)P(x, M)$ if $|N| \geq |M|$.

[6] DellaVigna, S. and U. Malmendier (2004): "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics* **119**, 353-402.

[7] Eliaz, K. and R. Spiegler (2006): "Contracting with Diversely Naive Agents," *Review of Economic Studies* **73**, 689-714.

[8] Eliaz, K. and R. Spiegler (in press): "Consumer Optimism and Price Discrimination," *Theoretical Economics.*

[9] Gabaix, X. and D. Laibson (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics* **121**, 505-540.

[10] Goldberg, L. (1993): "The structure of phenotypic personality traits," *American Psychologist* **48**, 26-34.

[11] Hoyer, Wayne D. (1984). "An Examination of Consumer Decision Making for a Common Repeat Purchase Product," *Journal of Consumer Research*, 11, 822-829.

[12] Manzini, P. and M. Mariotti (2007): "Sequentially Rationalizable Choice," *American Economic Review,* **97**, 1824-1839.

[13] Masatlioglu, Y. and E. Ok (2005): "Rational Choice with Status Quo Bias," *Journal of Economic Theory* **121**, 1-29.

[14] Masatlioglu, Y. and D. Nakajima (2007): "Choice by Constraint Elimination," Mimeo.

[15] Mullainathan S., J. Schwartzstein and A. Shleifer (2008): "Coarse Thinking and Persuasion," *Quarterly Journal of Economics* **123**, 577–619.

[16] Nedungadi, P. (1990): "Recall and Consumer Consideration Sets: Influencing Choice without Altering Brand Evaluations," *Journal of Consumer Research*, 17, 263-276.

[17] Piccione, M. and A. Rubinstein (2003): "Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns," *Journal of European Economic Association* **1**, 212-223.

[18] Roberts, J.H. and J.M. Lattin (1997): "Consideration: Review of Research and Prospects for Future Insights," *Journal of Marketing Research* **34**, 406-410.

[19] Rubinstein, A. (1988): "Similarity and Decision-Making Under Risk," *Journal of Economic Theory* **46**, 145-153.

[20] Rubinstein, A. (1993): "On Price Recognition and Computational Complexity in a Monopolistic Model," *Journal of Political Economy* **101**, 473-484.

[21] Rubinstein, A. and Y. Salant (in press): "(A,f): Choices with Frames," *Review of Economic Studies*.

[22] Shum, M. (2004): "Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast-Cereals Market," *Journal of Economics and Management Strategy*, 13, 241-272.

[23] Shapiro, J. (2006): "A 'Memory-Jamming' Theory of Advertising," mimeo, University of Chicago.

[24] Spiegler, R. (2005): "The Market for Quacks," *Review of Economic Studies* **73**, 1113-1131.

[25] Spiegler R. (2006): "Competition over Agents with Boundedly Rational Expectations," *Theoretical Economics* **1**, 207-231.

[26] Warner, M. (2006): "Salads or No, Cheap Burgers Revive McDonald's," *New York Times*, april 19.

[27] Varian, H. (1980): "A model of sales," *American Economic Review* **70**, 651-659.

[28] Zhang, S. and A.B. Markman (1998): "Overcoming the Early Entrant Advantage: The Role of Alignable and Nonalignable Differences," *Journal of Marketing Research*, 35, 413-426.