Competitive Framing*

Ran Spiegler†

February 20, 2013

Abstract

I present a simple framework for modeling two-firm market competition, where each firm chooses a market alternative as well as a costless "marketing message". Consumers follow a frame-dependent choice function, where the frame is a probabilistic function of the firms' marketing messages. This framework embeds several recent models in the literature on markets with boundedly rational consumers. I identify a property that consumer choice may satisfy, which extends the concept of Weighted Regularity due to Piccione and Spiegler (2012), and provide a characterization of Nash equilibria under this property. I use this result to analyze the equilibrium interplay between competition and framing in a large variety of models that fall into the framework.

1 Introduction

Boundedly rational consumers often make choices that are sensitive to the "framing" of individual alternatives or the choice set as a whole. For instance, the measurement units in which prices and other quantities are expressed may affect similarity judgments that guide consumers; variable font size in a contract may divert consumers' attention from one product attribute to another; the terms used to describe a lottery may induce consumers to categorize outcomes as gains or losses relative to a reference point and thus manipulate their risk preferences; and so forth.

*This project enjoyed financial support from the European Research Council, Grant no. 230251. I thank Beniamin Bachi, Elchanan Ben-Porath, Kfir Eliaz, Erik Eyster, Sergiu Hart, Ariel Rubinstein, Eilon Solan, seminar audiences at Chicago, Tel Aviv and Toronto and two anonymous referees for helpful comments.

†Tel Aviv University and University College London. URL: http://www.tau.ac.il/~rani. E-mail: r.spiegler@ucl.ac.uk.

Experimental psychologists have spent considerable effort and ingenuity eliciting such framing effects (see Kahneman and Tversky (2000)). Yet, there is a crucial difference between the psychologist's lab and the market setting in which frame-sensitive consumers are placed. While the psychologist's objective is to elicit the framing effect in order to expose decision makers' underlying choice procedures, the firm's objective is to maximize its profits. In a competitive market environment, it is not clear a priori how competitive pressures will affect the firms' incentive to manipulate consumer choice through framing. Will competition discipline or rather exacerbate framing effects? How does consumers' sensitivity to frames affect the competitiveness of the market outcome?

This paper presents a theoretical framework for exploring the interplay between framing and competition in a "Bertrand-like" market setting, in which two firms compete for one consumer according to a simultaneous-move game with complete information. Each firm i chooses a pair (a_i, m_i), where a_i is an "alternative" and m_i is a "marketing message". The set of feasible marketing messages may vary with the alternative offered. Selling one unit of an alternative a generates a net profit of $p(a)$ for the firm. Marketing messages are costless and enter the firms' payoff function only via their impact on consumer choice, which follows a two-step description. First, the profile of marketing messages (m_1, m_2) induces a "frame" f with probability $\pi_f(m_1, m_2)$. Second, the probability that the consumer chooses firm i is $s_i(a_1, a_2, f)$, where $s_1(a, b, f) = s_2(b, a, f)$ (firms' labels are irrelevant), and $s_1+s_2 \equiv 1$ (the consumer has no outside option). Thus, the frame f is simply a parameter in the consumer's probabilistic choice function.

The interpretation of the distinction between "alternative" and "marketing message" is that the former consists of intrinsic, utility-relevant features of the firm's offer, whereas the latter consists of utility-irrelevant details of its description. We will see that it is not always clear where to draw the line between the two, and occasionally this modeling decision will be dictated by analytic convenience. In particular, in the absence of any structure on the primitives of the modeling framework, it can incorporate any probabilistic choice behavior (including standard random-utility maximization), as long as it satisfies the label-neutrality and no-outside-option restrictions. The interest in this modeling framework is that it suggests restrictions that *do* constrain its generality.

The modeling framework can capture a variety of market situations in which firms manipulate consumer choices: shrouding product attributes (Gabaix and Laibson (2006)), obfuscation by means of price variation across numerous dimensions (Spiegler (2006)), or using incommensurable measurement units to reduce comparability of market alternatives (Piccione and Spiegler (2012), PS henceforth). This paper will introduce a number of new examples that fall into the modeling framework and capture a variety of framing effects in competitive marketing settings, such as narrow bracketing of financial risk, or using product classification to influence perceived product substitutability. The following detailed example previews one of these applications.

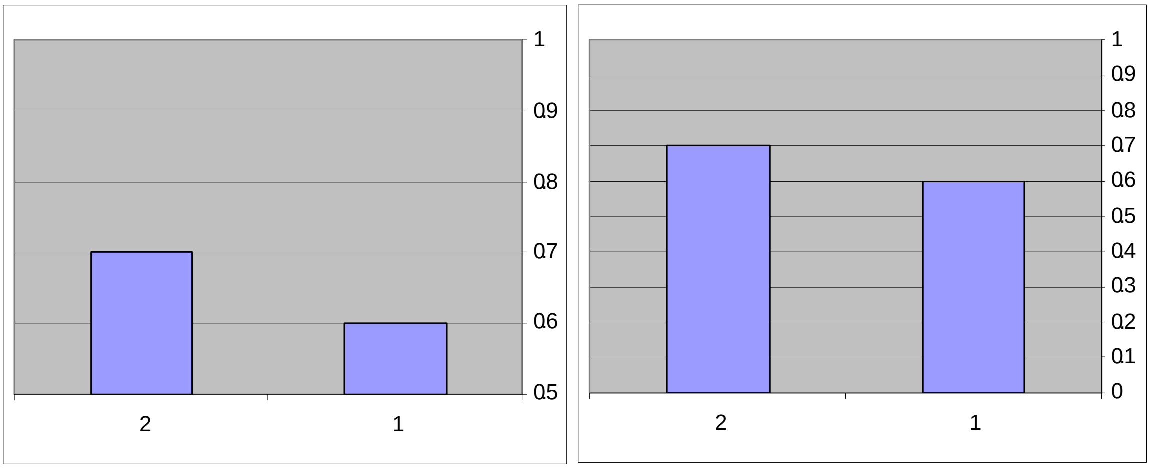

(Figure 1: Manipulation of similarity judgments($f = 0.5$ on the left, $f = 0$ on the right))

Example 1.1: Manipulating the scale of similarity judgments
Identify a market alternative with its price $p \in [0, 1]$. Consumer choice is based on similarity judgments that can be manipulated by framing. Imagine that prices are represented by a column chart, as in Figure 1. A given price difference may appear large or small, depending on the chart's scale. A consumer with a "similarity coefficient" α chooses the cheaper firm whenever $|p_2 - p_1| / (1 - f) > \alpha$. The frame f is thus the "origin" of the graphical representation the consumer relies on to form intuitive similarity judgments. Assume that α is distributed according to some *cdf* G, such that if $p_1 < p_2$, firm 1's market share under the frame f is $\frac{1}{2}\{1 + G[(p_2 - p_1)/(1 - f)]\}$.

How can firms manipulate the consumer's frame? The set of feasible marketing messages given p is $(-\infty, p)$. The interpretation is that each firm i uses its own column chart to present its price p_i. The scale of the chart's vertical axis is $1 - m_i$. The number m_i is thus the origin of the firm's chart. The consumer constructs his own graphical representation, taking the firms' graphs as inputs. When $m_1 = m_2 = m$, he automatically adopts this origin and it becomes his frame - that is, $f = m$. In contrast, when $m_1 \neq m_2$, the inconsistency between the coordinate systems of the two firms' graphs alerts the consumer to the framing effect. In response, he reverts to his own "default frame" $f = 0$.

My analysis of virtually all the examples in the paper is unified by a property that consumers' choice function may satisfy. We will say that π satisfies "*Weighted Regularity*" (**WR**) if there exist a mixture λ over $\cap_{a\in A}M(a)$ and a distribution v over F, such that if one firm randomizes over marketing messages according to λ, the distribution over F is σ, independently of the other firm's marketing. Weighted regularity thus means that consumer choice is potentially neutral to marketing. This concept was originally introduced by PS in the context of a model which is a special case of the current framework. I reformulate and extend the concept to fit the more general environment. Weighted regularity itself is a special case of a more general property of (s, π), called "*Worst-Case Independence of Marketing*" (**WIM**), which means that for every alternative a that a firm may offer, there is a mixture λ^a over the set of feasible marketing messages that max-minimizes the firm's expected market share, independently of the alternative offered by the rival firm.

The basic result of this paper is that under WIM, firms choose alternatives in Nash equilibrium as if their marketing strategy for any a is λ^a (subject to equilibrium existence considerations). This result is an elementary application of the Minimax Theorem. This fundamental theorem is applicable because by the no-outside-option property of s, the firms' choices of marketing messages for every profile of alternatives (a_1, a_2) constitute a zero-sum game. The basic result means that in equilibrium, firms' market shares for any realization (a_1, a_2) are determined by the value of the associated zero-sum game. When π satisfies WR, this automatically implies WIM, and the basic result implies that firms choose alternatives as if the distribution over the consumer's frame is exogenously given by σ as defined above.

WIM and WR are not "generic" properties (although in some cases they may emerge from a larger model that endogenizes s). Nevertheless, it turns out to hold in surprisingly many examples (possibly because examples, by their nature, tend to impose "non-generic" symmetries), and thus facilitates their analysis. In some cases, observing that WR holds enables genuine insights into the equilibrium interplay between competitive forces and framing effects. For illustration, consider Example 1.1, and observe that each firm can unilaterally enforce the consumer's default frame $f = 0$ by playing $m = 0$. Hence, WR is trivially satisfied. By the basic result, firms choose prices in Nash equilibrium as if the consumer always adopts the default frame. This in turn leads to the following result. If $G \equiv U[0, k]$ for some $k < 1$, the game has a unique Nash equilibrium, in which both firms play $p = k$ and $m = 0$. That is, firms totally refrain from trying to manipulate the consumer's frame. In this sense, competition disciplines framing (a monopolist facing a consumer with some outside option would strictly prefer to manipulate the consumer's frame). Yet the equilibrium outcome is not more competitive for that. Indeed, if a regulator could enforce a frame somewhere between 0 and $k/2$, consumers would become more sensitive to small price differences and the equilibrium outcome would be more competitive. The lesson from this example is that even when competition puts a break on firms' attempt to manipulate consumer choice, this does not necessarily mean that the market outcome is as competitive as it could be, even when the constraints on consumer rationality are taken into account. The other applications in the paper similarly provide insights into framing in competitive environments.

Related literature

This paper belongs to the literature on "Behavioral Industrial Organization", reviewed by Ellison (2006) and Armstrong (2008), and synthesized into a textbook presentation in Spiegler (2011). One of the earliest papers in this literature, Rubinstein (1993), was also the first to incorporate framing into a model of optimal pricing. Specifically, Rubinstein examined a monopolist who frames a price p by splitting it into two components, p^1 and p^2, such that $p^1 + p^2 = p$, and commits ex-ante to a state-contingent probability distribution over price vectors. Consumers are limited in their ability to compute the actual price from a given vector. Specifically, they are restricted to functions that can be computed by an array of low-order perceptrons, and they choose the actual function after the firm has committed to its strategy and before the state is realized. Rubinstein assumes that consumers differ in the order of their perceptrons as well as in the state-contingent cost of providing the product to them. When these two traits are negatively correlated, the monopolist has a strategy that effectively screens the consumer's type and approximates the first-best.

So far, the behavioral I.O. literature has progressed by considering specific choice models that capture particular aspects of consumer psychology. A natural response to this proliferation of examples is to raise the level of abstraction and seek common features that override the specific psychological phenomena such models incorporate. Eliaz and Spiegler (2011) and PS are steps in this general direction. Both papers develop market models in which consumers' propensity to make preference comparisons is a general function of utility-relevant features of market alternatives (such as price or quality) as well as utility-irrelevant features (i.e., their "framing"). From a narrowly technical point of view, the present paper makes a modest contribution, as the basic result described above reformulates and extends part of Theorem 1 in PS. The main contribution of the present paper is to show that by extending the Piccione-Spiegler formalism to a much wider class of frame-sensitive choice procedures, we can analyze market implications of many phenomena of consumer psychology, and classify them according to their impact on the competitiveness of market equilibrium.

The paper is also related to the recent decision-theoretic literature on "choices with frames". Masatlioglu and Ok (2005) were to my knowledge the first to axiomatize choice correspondences defined over "extended choice problems" that also specify whether one of the feasible alternatives is a status quo. Salant and Rubinstein (2008) and Bernheim and Rangel (2009) systematized and generalized the notion of extended choice problems with frames and used it to analyze questions like the rationalizability of frame-dependent choice functions, or the elicitation of welfare from observed frame-sensitive choices. The function s in the present model is a probabilistic extension of this concept of choices with frames. The main difference is that the frame is endogenously determined by the firms' marketing strategies.

2 A Modeling Framework

A market consists of two firms and a single consumer. The firms play a symmetric simultaneous-move game with complete information, which is based on the following primitives:

- A set A of *alternatives*.
- *For every* $a \in A$, a set of feasible *marketing messages* $M(a)$. Define $M = \cup_{a\in A} M(a)$.
- A set F of *frames*.
- A symmetric function $\pi : M \times M \rightarrow \Delta(F)$, such that $\pi_f(m_1, m_2)$ is the probability that the consumer adopts the frame $f \in F$ when the firms' profile of marketing messages is (m_1, m_2).
- A function $p : A \rightarrow \mathbb{R}$ that specifies the net profit to the firm from selling one unit of any alternative.
- A probabilistic choice function s, where $s_i(a_1, a_2, f) \in [0, 1]$ is the probability that the consumer chooses firm i given that the profile of alternatives is (a_1, a_2) and the consumer has adopted the frame f. I assume that $s_1 + s_2 \equiv 1$ and

$s_1(a, b, f) = s_2(b, a, f)$ for every $a, b \in A$. That is, the consumer is forced to choose one of the firm, and his choice is not sensitive to the firms' labels.

For expositional simplicity, the general description of the model will take the sets A, M, F to be finite. Note, however, that in many applications, some of these sets are infinite; extending the formal model and the basic result of Section 3 to this case is straightforward.

A pure strategy for a firm is a pair (a, m), where $a \in A$ and $m \in M(a)$. Each firm $i \in \{1, 2\}$ chooses (a_i, m_i) to maximize

$$p(a_i) \cdot \left[\sum_f \pi_f(m_1, m_2, f) \cdot s_i(a_1, a_2, f)\right]$$

I will sometimes use

$$s_i^*((a_j, m_j)_{j=1,2}) = \sum_f \pi_f(m_1, m_2, f) \cdot s_i(a_1, a_2, f)$$

to denote firm i's expected market share given the strategy profile $(a_j, m_j)_{j=1,2}$. I refer to s and s^* as the choice function and market share functions, respectively.

When no restrictions are imposed on the primitives, the modeling framework is behaviorally empty, in the sense that it can accommodate any probabilistic choice function that satisfies the symmetry and no-outside-option assumptions. To see why, note that we could always set $M = A$, $M(a) = \{a\}$, $F = \{f\}$. In particular, conventionally rational choice , described by maximization of some random utility function over $A \times M$, is not ruled out by the model. The value of this elaborate framework lies in the language it provides for unifying behavioral I.O. models, as well as in the fruitful restrictions on choice behavior that it suggests.

The following are examples of models that fit into this framework.

Example 2.1: Shrouded attributes (a variant on Gabaix and Laibson (2006))
An alternative is a vector $(a^1, a^2) \in [0, \infty)^2$ specifying the quality level of two product attributes, where $p(a^1, a^2) = 1 - (a^1 + a^2)$. The set of feasible marketing messages is $\{0, 1\}$, regardless of the offered alternative. The interpretation of $m = 1$ is that the firm "shrouds" the second attribute. The set of frames is also $\{0, 1\}$, where $f = 1$ means that the second attribute ends up being shrouded. Let $\pi_1(m_1, m_2) = 1 - \lambda + \lambda m_1 m_2$ for all (m_1, m_2). The consumer's choice function s selects the firm i that maximizes $a_i^1 + (1 - f) \cdot a_i^2$, with a symmetric tie-breaking rule. The interpretation is that if

both firms shroud the second attribute, the consumer is unaware of it and chooses entirely according to the first attribute. If at least one firm "unshrouds" the second attribute, then with probability λ the consumer becomes "enlightened" and aggregates both attributes.

Example 2.2: Limited comparability of description formats (PS)
An alternative is identified with the product price $p \in [0,1]$. The set of marketing messages is a finite set M, independently of the price that firms charge. An element in M is interpreted as a "description format" (e.g., a measurement unit in which the price is stated). The set of frames is $\{0,1\}$, where $f = 1$ means that the firms' chosen formats are comparable. Assume $\pi(m,n) \equiv \pi(n,m)$. The choice function is $s_1(p_1,p_2,f) = \frac{1}{2}[1 + f \cdot sign(p_2 - p_1)]$. The interpretation is that when the firms' formats are comparable, the consumer is able to make a price comparison and correctly selects the cheaper firm (with a symmetric tie-breaking rule). When he cannot make a comparison, he chooses each firm with probability $\frac{1}{2}$.[1]

Example 2.3: Obfuscation as noise (Spiegler (2006))
Firms sell a homogenous product. Each firm simultaneously chooses a probability measure over its price that may take values in $(-\infty, 1]$. The consumer draws one sample point from each distribution and chooses the cheapest firm in his sample (with symmetric tie breaking). The actual price he pays, however, is the mean of the chosen firm's distribution. To cast this model into the current framework, let $A = (-\infty, 1]$ be the set of possible prices. Let $M(a)$ be the set of all zero-mean probability measures over $(-\infty, 1-a]$. Let F be the set of real numbers, and π is defined as a convolution:

$$\pi_f(m_1, m_2) = \int dm_1(x) dm_2(x+f) dx$$

Assume that s selects firm 1 (2) whenever $a_1 < a_2 + f$ ($a_1 > a_2 + f$). Tie-breaking is symmetric. Thus, a frame f represents a bias against firm 2 that takes the form of a tax on its price, which is equal to the difference between the firms' (independent) individual noise realizations.

When firms play a Nash equilibrium in our game, the induced profile of (possibly mixed) marketing strategies is formally equivalent to Nash equilibrium in a Bayesian *zero-sum* game, where: (*i*) a_i corresponds to player i's type, such that a state consists of the pair (a_1, a_2); (*ii*) player i's (type-dependent) action set is $M(a_i)$; (*iii*) player i's

[1] Carlin (2009) and Chioveanu and Zhou (2011) analyze special cases of π and extend the choice model to $n > 2$ firms.

payoff function is s_i^* (m_1 and m_2 are the players' actions, while a_1 and a_2 serve as parameters); and (*iv*) the marginal probability distribution over A induced by firm i's equilibrium strategy corresponds to a prior distribution over its type. Because Nash equilibrium strategies are by definition statistically independent, this means that the players' types in the analogous Bayesian game are independently distributed. I will assume that for every (a_1, a_2), the auxiliary "ex-post" zero-sum game $< M(a_1), M(a_2), s_1^* >$ has a value, denoted $v(a_1, a_2)$. Of course, the existence of a value is guaranteed when M is finite. We will draw on this formal analogy in the sequel.

Discussion: What is a frame?
Following Salant and Rubinstein (2008) and Bernheim and Rangel (2009), I model a frame as a parameter in the consumer's (probabilistic) choice function. Associating a frame with the entire choice set is natural in many cases, e.g. when the frame is a reference point such as a status quo, or a scale as in Example 1.1. In other cases, it would be more natural to associate frames with the individual alternatives rather than with the choice set as a whole (this was the approach taken in Eliaz and Spiegler (2011) and PS), and the latter is a formal contrivance that serves our analytic purposes.

The distinction between "alternatives" and "marketing messages" is not always clear-cut. For instance, in Example 2.1, one could argue that splitting a product into two attributes is part of its marketing, rather than being intrinsic to the product. Similarly, the packaging of products often has functional aspects that are also useful for marketing purposes (e.g., an unusual shape attracts attention). In these cases, the distinction between a and m is somewhat arbitrary. As we will see later in the paper, I will occasionally draw the demarcating line according to considerations of analytic convenience, overriding a more natural distinction.

3 Worst-Case Independence of Marketing

This section identifies a property that the market share function s^* may satisfy, and which turns out to hold in a number of applications. I then show its role in facilitating Nash equilibrium analysis.

Definition 1 *A market share function s^* satisfies **Worst-Case Independence of Marketing (WIM)** if for every $a_1 \in A$ there exists $\lambda_1^* \in \Delta(M(a_1))$ that solves the* *max-minimization problem*

$$\max_{\lambda_1 \in \Delta(M(a_1))} \min_{\lambda_2 \in \Delta(M(a_2))} \sum_{m_1} \sum_{m_2} \lambda_1(m_1)\lambda_2(m_2) s_1^*((a_j, m_j)_{j=1,2})$$

*for all $a_2 \in A$. We will say that s^* satisfies **Strong WIM** if λ_1^* is constant over A.*

WIM means that a firm's worst-case-optimal marketing strategy is independent of the alternative that its opponent offers. Strong WIM means that it is also independent of the firm's own choice of alternative. WIM is not meant to be justifiable on a priori grounds, nor is it meaningful from a single-agent behavioral perspective. Moreover, it is a rather superficial property, because it does not exploit the finer structure of the function s^* that involves the primitives π and s. Nevertheless, when the primitives have some structure, WIM often has a natural behavioral interpretation, or simply holds automatically.

The following is a special case of WIM that will be of special interest in the sequel.

Definition 2 *The function π satisfies **Weighted Regularity (WR)** if there exist $\lambda \in \Delta(\cap_{a \in A} M(a))$ and $\sigma \in \Delta(F)$, such that*

$$\sum_m \lambda(m)\pi(m, n) = \sigma$$

for every $m \in \cup_{a \in A} M(a)$. In this case we say that WR is verified by (λ, σ).

This property was originally introduced by PS in the context of the model given by Example 2.2, where it means that each firm can unilaterally enforce a constant comparison probability. Its general interpretation is that firms are potentially neutral to marketing, in the sense that one firm can always (that is, regardless of the alternative it offers) enforce a distribution over the consumer's adopted frame that is independent of its opponent's marketing strategy.

Lemma 1 *If π satisfies WR, then s^* satisfies Strong WIM.*

Proof. Consider the zero-sum game induced by (a_1, a_2). Suppose that firm 1 mixes over $M(a_1)$ according to the strategy λ as in the definition of WR. Then, since this enforces a distribution over f that is independent of firm 2's marketing message, firm

2 is indifferent among all marketing messages, hence (λ, λ) is a Nash equilibrium in the zero-sum game. By the Minimax Theorem, this means that λ max-minimizes firm 1's expected market share for all (a_1, a_2), hence Strong WIM holds. ■

Both Examples 1.1 and 2.1 trivially satisfy WR, because each firm can unilaterally enforce the deterministic frame $f = 0$. On the other hand, Example 2.3 necessarily violates WIM. Observe that the ex-post zero-sum game associated with (a_1, a_2) in this example is what Hart (2008) referred to as a continuous General Lotto game $\Lambda(1-a_1, 1-a_2)$. According to Theorem 1 in Hart (2008), there is a unique equilibrium in this game, where player 1's max-min strategy varies with a_2.

Comment: WIM and rational choice
As mentioned above, the modeling framework does not rule out conventionally rational choice. The following is an example. Let $A = \mathbb{R}_+$ be the set of feasible product prices, and let $M(a) = M$ represent a set of product types. Define $F = M^2$, and assume that $\pi(m_1, m_2)$ assigns probability one to (m_1, m_2). The choice function s selects the firm i that maximizes $u(m_i) - a_i$ (with a symmetric tie-breaking rule), where $(u(m))_{m \in M}$ is drawn from some probability measure over $\mathbb{R}_+^K$.

Is there a logical link between this form of rational choice and WIM? Suppose that for every $m \in M$, $u(m)$ is independently drawn from some *cdf* G_m, such that the *cdf*s can be ordered by FOSD. Then, Strong WIM is trivially satisfied, because firm 1 can maximize its market share for any p_1, p_2, m_2 by selecting the marketing message associated with a *cdf* that FOSD the *cdf*s associated with all other marketing messages. However, it is easy to generate examples that violate WIM. For instance, let $M = \{m, n\}$, and suppose that $(u(m), u(n))$ is $(1, 0)$ with probability $\frac{2}{3}$ and $(0, 1)$ with probability $\frac{1}{3}$. Fix p_1. When $p_2 = p_1$, there is a unique Nash equilibrium in the zero-sum game, where both firms play m. However, when $0 < p_2 < p_1 < 1$, this is no longer an equilibrium, because firm 1 has an incentive to deviate to n. Therefore, there exists no max-minimizing strategy for firm 1 that is independent of firm 2's price.

The following simple result will serve us in the applications section. It shows that under WIM, Nash equilibrium (if one exists) has the property that the firms' choices of alternatives are taken as if they expect they become commonly known before firms take their marketing decisions. The proof is an elementary application of the Minimax Theorem, and it extends an idea that is "buried" in the proof of Theorem 1 in PS. I carry the expositionally convenient assumption that A, M, F are all finite sets, such that existence of mixed-strategy Nash equilibrium is ensured.

Proposition 1 *Suppose that s^* satisfies WIM. Then, in Nash equilibrium, firm 1's expected market share conditional on (a_1, a_2) is $v(a_1, a_2)$ for almost every (a_1, a_2).*

Proof. Consider a Nash equilibrium, and let μ_i denote the marginal probability distribution over A induced by firm i's equilibrium strategy. Fix an alternative $a_1 \in A$. By assumption, the zero-sum game associated with every (a_1, a_2) has a value $v(a_1, a_2)$. By WIM, there exists a feasible mixed marketing strategy that max-minimizes firm 1's expected market share, independently of a_2. By the Minimax Theorem, firm 1 can ensure an expected market share of at least $v(a_1, a_2)$ for all a_2. Let S_1^* denote firm 1's ex-ante expected market share in equilibrium. Then,

$$S_1^* \geq \sum_{a_1} \sum_{a_2} v(a_1, a_2)\mu_2(a_2)\mu_1(a_1)$$

By definition of the value, $v(a_2, a_1) = -v(a_1, a_2)$. Therefore, by a similar argument, for every a_2, firm 2 can play a marketing strategy that ensures an expected market share of at least $1 - v(a_1, a_2)$ for all a_1. Then,

$$S_2^* \geq \sum_{a_2} \sum_{a_1} (1 - v(a_1, a_2))\mu_1(a_1)\mu_2(a_2)$$

such that $S_1^* + S_2^* \geq 1$. By assumption, $S_1^* + S_2^* = 1$. Therefore, the above inequalities must both be binding, which means that firm 1's expected market share conditional on (a_1, a_2) is equal to $v(a_1, a_2)$ for a probability-one set of realizations (a_1, a_2). ■

Thus, when the market share function satisfies WIM, Nash equilibrium analysis can be reduced to a "backward induction" algorithm, which is often easy to carry out. We first obtain the market shares conditional on the offered alternatives, by finding the value of the zero-sum game associated with every profile of alternatives. We then plug these market shares into the firms' payoff function as if they only choose alternatives. Thus, equilibrium pricing behavior is the same as in a two-stage variant on the model, in which firms choose alternatives in the first stage and marketing messages in the second stage, after observing the realization of their first-stage strategy. WIM thus ensures that equilibrium predictions are "robust" to this modification of the game's move sequence. Put more crudely, it does not matter whether firms choose a "veneer" for their product after, or simultaneously with their choice of its utility-relevant features.[2]

[2] The reason I insist on the simultaneity of choices of alternatives and marketing messages to begin with, is that I want to enable firms to consider all aspects of their competitive strategy - product design, pricing and marketing - as an integrated decision. This fits many real-life situations For instance, if

Corollary 1 *Suppose that WR is verified by* (λ, σ)*. Then, in Nash equilibrium,*

$$s_1^*(a_1, a_2) = \sum_{f \in F} \sigma_f \cdot s_1(a_1, a_2, f)$$

for almost every (a_1, a_2)*.*

Proof. Since π satisfies WR, Lemma 1 implies that s^* satisfies Strong WIM. By Proposition 1, in Nash equilibrium, market shares for every (a_1, a_2) is given by the value of the associated zero-sum game. As we observed in the proof of Lemma 1, λ max-minimizes firm 1's market share for almost every (a_1, a_2). By definition, λ induces the distribution σ over F, independently of the other firm's marketing message. It follows that for almost every (a_1, a_2),

$$s_1^*(a_1, a_2) = v(a_1, a_2) = \sum_{f \in F} \sigma_f \cdot s_1(a_1, a_2, f)$$

■

Thus, under WR, firms' equilibrium choices of alternatives are made as if the distribution over the consumer's frame is exogenously given by σ.

4 Applications

In this section I analyze several market models with boundedly rational consumers. In all cases, applying Proposition 1 greatly simplifies the equilibrium analysis. Each application is meant to convey three types of messages: (*i*) *modeling*: showing how the framework can capture an aspect of competitive framing that has not hitherto been addressed in the economic literature; (*ii*) *economic*: providing a lesson about the interplay between competitive forces and framing in specific market situations; (*iii*) *methodological*: demonstrating an aspect of WR or WIM that may appear in other applications. I summarize these lessons at the end of each sub-section.

4.1 Default Frames

When experimental psychologists attempt to elicit a framing effect, they use an inter-subject methodology, such that different groups are exposed to different frames. The

one supermarket chain revises its expectation of a competing chain's strategy, it will reconsider both its pricing strategy and how to advertise it.

reason is that in many cases, a framing effect is like a magician's trick, and exposure to conflicting frames is like knowing the inner workings of the trick: it ruins the illusion and annuls the effect. For example, the famous "Asian flu" experiment (Kahneman and Tversky (1979)), which demonstrated the sensitivity of risk attitudes to the description of outcomes in terms of gains or losses, relied on this technique: if subjects had been exposed to both treatments at the same time, most of them would probably have seen through the trick and switched to a more consequentialist mode of thinking. Unlike experimental psychologists, firms in our model are not interested in eliciting a framing effect per se; they will exploit consumers' sensitivity to frames only if this serves the profit-maximization objective. The question is whether competitive forces rule out the elicitation of framing effects, when these are annulled by the presence of conflicting frames. This sub-section explores this question.

Consider a specification in which A is a compact and convex subset of some Euclidean space. Let $M(a) = F$ for all $a \in A$, where F is finite. One of the elements in F, denoted 0, is referred to as the "default frame". Assume that whenever $m_1 = m_2 = f$, π assigns probability one to f. In contrast, whenever $m_1 \neq m_2$, π assigns probability one to the default frame. Let us begin with a result that imposes very little structure on the consumer's frame-sensitive choice function. Specifically, assume that whenever $a_1 \neq a_2$, $f \neq g$ implies $s(a_1, a_2, f) \neq s(a_1, a_2, g)$. Note that I do not require the default frame to induce "more rational" choices than the other frames.

Proposition 2 *In any symmetric Nash equilibrium in which firms' marginal distribution over A is atomless, firms play $f = 0$ with probability one.*

Proof. Each firm can unilaterally induce the frame $f = 0$, by playing $m = 0$. Therefore, WR trivially holds. By Corollary 1, firm 1's market share is $s_1(a_1, a_2, 0)$ for almost every (a_1, a_2). If a frame $f \neq 0$ is played with positive probability in equilibrium, then by the assumption that the firms' equilibrium strategies induce an atomless marginal distribution over A, there must be a positive-measure set of realized action profiles for which $a_1 \neq a_2$ and $m_1 = m_2 = f$. For a profile $((a_1, f), (a_2, f))$ in this set, firm 1's market share is $s_1(a_1, a_2, f)$, which is different from $s_1(a_1, a_2, 0)$ by assumption, a contradiction. ■

Thus, symmetric Nash equilibrium in this game selects the default frame, as long as the marginal equilibrium distribution over alternatives is atomless. The logic behind the result is somewhat reminiscent of "no trade" theorems. We have already commented on the analogy between the current modeling framework and Bayesian zero-sum games. Models of speculative bilateral trade with asymmetrically informed traders fall into this class of games. In the current model, each firm can unilaterally enforce the default frame, just as in models of speculative trade, each trader can unilaterally enforce the no-trade outcome by refusing to trade. In both models, equilibrium selects this unilateral action, as a result of similar strategic considerations.

Of course, this result is limited because it is stated for a certain class of equilibria that need not exist. Let us now examine examples that impose more structure on s. Consider Example 2.1. We noted that WR holds because each firm can unilaterally enforce $f = 0$. By Corollary 1, firms choose (a^1, a^2) as if consumers choose rationally (i.e., $f = 0$). Therefore, in Nash equilibrium, both firms play $a^1 + a^2 = 1$ with probability one. If firm 1, say, assigns positive probability to $m_1 = 1$, then firm 2 can profitably deviate to a $m_2 = 1$ coupled with mixing between $(a^1, a^2) = (0, 1-\varepsilon)$ and $(a^1, a^2) = (1-\varepsilon, 0)$, where $\varepsilon > 0$ is arbitrarily small. Thus, in Nash equilibrium it must be the case that $m_1 = m_2 = 0$. In other words, both firms unshroud the second attribute.

Let us now revisit Example 1.1, which describes a concrete story in which moving the consumer's frame away from the default requires that both firms coordinate on the same frame. In this case, the marginal equilibrium distribution over alternatives is not atomless, hence Proposition 2 does not apply. Let us prove the result that in Nash equilibrium, firms play $(k, 0)$. First, by Proposition 1, firms choose prices as if firm 1's expected market share from every (p_1, p_2) is $s_1^*(p_1, p_2, 0)$ - that is, as if they play a game in which each firm i chooses $p_i \in [0, 1]$ to maximize $p_i \cdot \frac{1}{2}[1 + G(p_j - p_i)]$. Now impose the assumption that $G \equiv U[0, k]$. It is easy to show that this auxiliary game has a unique Nash equilibrium, in which $p_1 = p_2 = k$. It remains to verify that firms play $m = 0$ with probability one. It is clear that this profile of marketing messages is consistent with equilibrium, because neither firm can unilaterally change the consumer's frame when $m_1 = m_2 = 0$. Let us show that no other equilibrium exists. Suppose that firm 2, say, plays some mixture $\lambda \in \Delta(F)$ over marketing messages, which assigns positive probability to some $f \neq 0$. Suppose that firm 1 deviates to the pure strategy $(k - \varepsilon \cdot sign(f), f)$, where $\varepsilon > 0$. Note that whenever firm 2's realized marketing message is equal to (different from) f, the induced frame is f (0). It is now straightforward to calculate that if ε is sufficiently small, firm 1's deviation is profitable.

This example conveys a subtle point regarding the interplay between competition and framing effects. The zero-sum aspect of the competitive interaction, coupled with the feature that each firm can unilaterally enforce the default frame, implies that firms refrain from manipulating the consumer's frame in equilibrium. In this sense, competition disciplines obfuscation in this model (unlike Spiegler (2006), where competition in fact exacerbates the firms' obfuscatory strategy in equilibrium). In contrast, a monopolistic firm facing a consumer who has some exogenous outside option would actively engage in framing: if it chooses to offer a more (less) attractive alternative than the outside option, it will use framing to exaggerate (downplay) the difference.

However, the fact that competition disciplines framing does not mean that it makes the market outcome more favorable to consumers. Indeed, if firms coordinated on a frame $f \in (0, \frac{k}{2})$, possibly under a regulator's instruction, the equilibrium price would be $k(1-f)$, hence the market outcome would be more competitive. The intuition is that setting f above the default level enhances the consumer's sensitivity to small price differences, and this strengthens competitive pressures. Thus, the mere finding that competition eliminates obfuscation does not imply that the outcome is more competitive than if consumer choices were manipulated.

The lessons from this sub-section can be summarized as follows: (i) the modeling framework can accommodate market situations in which consumers may face conflicting frames suggested or implied by the competing firms, and thus forces the modeler to make an explicit assumption about the way consumers respond to this tension; (ii) if consumers respond to such a conflict by adopting a "default" frame, then in equilibrium firms never try to manipulate the consumer's frame away from this default; (iii) when a firm can unilaterally enforce the consumer's frame, we have a simple case of WR.

4.2 Bracketing Financial Risk

In this sub-section I study a model in which firms compete in lotteries for a risk averse expected-utility maximizing consumer, where the carrier of the consumer's vNM utility function can be manipulated through framing. A well-known example for this phenomenon involves narrow bracketing of financial risk. It has been noted by experimentalists (see, for instance, Rabin and Weizsäcker (2009) and Eyster and Weizsäcker (2011)) that when investors face a monetary lottery that is correlated with another source of financial risk, they may treat the gains and losses defined by each lottery in isolation, or combine the two into one grand lottery over their wealth, depending on how the decision problem is framed.

Another example involves the description of interest rates in nominal or real terms. In the presence of inflation, this element of framing may affect consumers' perception of the riskiness of financial products. For instance, a deposit account bearing 2%-plus-inflation interest would appear safe when presented as such ("you are guaranteed 2%

in real terms"), yet risky when considered from a nominal point of view. Conversely, a deposit account bearing 5% nominal interest would appear safe when presented as such ("you get 5% no matter what"), yet a risky gamble when evaluated in real terms. Thus, firms competing in such products may strategically frame them in nominal or real terms in order to manipulate consumers' risk preferences. (Shafir, Diamond and Tversky (1997) provide experimental evidence for this effect.)

The model I analyze aims to capture this phenomenon. I use the interest-cum-inflation story, but I treat "inflation" additively rather than multiplicatively, for tractability. Assume that firms provide liquidity in return for interest payments. Let ε be a positive-valued random variable representing the rate of inflation, the mean value of which is μ. An alternative is a loan contract $a = (r, \alpha)$, where $r \in \mathbb{R}$ is the stated interest rate and $\alpha \in \{0, 1\}$ indicates whether it is indexed to inflation. The actual nominal interest rate induced by an offer a is a real-valued random variable $r + \alpha\varepsilon$. The actual real interest rate induced by a is $r - (1 - \alpha)\varepsilon$. I use $p = r - (1 - \alpha)\mu$ to denote the expected real interest rate, and assume that $p(a) = p$ - that is, the firm is risk neutral and its profit from a loan contract is the expected real interest it bears.

Assume that $M(r, \alpha) = \{\alpha\}$. Indexing interest to inflation ($\alpha = 1$) is interpreted as implying a "real frame", thus encouraging the consumer to think about outcomes in real terms that take inflation into account. In contrast, failure to index interest to inflation ($\alpha = 0$) implies a "nominal frame", thus encouraging the consumer to think about outcomes in nominal terms that ignore inflation. Let $F = \{0, 1\}$, where $f = 1$ (0) means that the consumer adopts a real (nominal) frame. I assume that π assigns probability $\frac{1}{2}$ to m_1 and m_2 each. (Thus, if $m_1 = m_2 = m$, then $f = m$.) The interpretation is that the consumer surveys the firms' offers sequentially in random order, and he adopts the frame implied by the first offer he considers.

To complete the model, we need to describe the consumer's choice function. Assume that he is endowed with a concave vNM utility u from money that exhibits CARA. However, the domain to which this function is applied is frame-sensitive. Specifically, the consumer chooses the firm i that maximizes

$$Eu[-(r_i + \alpha_i\varepsilon) + f\varepsilon]$$

with a symmetric tie-breaking rule. Let $c = |Eu(\varepsilon - \mu)|$ be the certainty equivalent of the random variable $\varepsilon - \mu$ representing unanticipated inflation.[3]

[3]Note there is no clear distinction between substance and framing in this model, because the aspect of the contract that affects the consumer's frame - i.e., whether it pegs interest to inflation - is also utility-relevant.

To illustrate the consumer's choice rule, suppose that he faces a choice between two contracts, $a_1 = (r, 1)$ and $a_2 = (r+\mu, 0)$, which are characterized by the same expected real interest. If the consumer adopts a real (nominal) frame, he views a_1 (a_2) as a sure thing and a_2 (a_1) as risky, hence he will choose a_1 (a_2). We can see that because framing manipulates the domain to which the consumer applies his risk preferences, it can cause preference reversals.

Proposition 3 *The game has a unique symmetric Nash equilibrium. Firms randomize over the expected real interest rate p according to the cdf*

$$G(p) = \frac{3}{2}(1 - \frac{c}{2p})$$

defined over the support $[\frac{c}{2}, \frac{3c}{2}]$, and independently randomize uniformly between $\alpha = 0$ and $\alpha = 1$. Equilibrium industry profits are $\frac{3c}{4}$.

Proof. The proof is based on a redefinition of the elements of the model (namely, A, M, F, π, s) such that WR can be applied. Given a contract (r, α), the alternative is identified with the expected real interest rate p implicit in the contract, whereas the marketing message is α, such that $M(p) = \{0, 1\}$ for all p. Redefine $F = \{0, -c, c\}$, and assume that $\pi = (\pi_0, \pi_{-c}, \pi_c)$ is given as follows:

$$\pi(m_1, m_2) = \begin{cases} (1, 0, 0) & if \quad m_1 = m_2 \\ (0, \frac{1}{2}, \frac{1}{2}) & if \quad m_1 \neq m_2 \end{cases}$$

Finally, the choice function is

$$s_1(p_1, p_2, f) = \frac{1}{2}[1 + sign(p_2 - p_1 - f)]$$

This formulation is equivalent to the original one in terms of the firms' payoff function.

Note that π satisfies WR, since it is verified by $\lambda(0) = \lambda(1) = \frac{1}{2}$ and $\sigma = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. Therefore, by Corollary 1, in Nash equilibrium the market share function is

$$s_1^*(p_1, p_2) = \begin{cases} 1 & if \quad p_2 - p_1 > c \\ \frac{7}{8} & if \quad p_2 - p_1 = c \\ \frac{3}{4} & if \quad p_2 - p_1 < c \end{cases} \qquad (1)$$

for all price profiles (p_1, p_2) for which $p_1 < p_2$ (and s_2^* is symmetrically defined).

The appendix establishes that this reduced price-competition game has a unique symmetric mixed-strategy Nash equilibrium given by G. As to the firms' indexation decisions, if they play $\alpha = 1$ with a probability different than $\frac{1}{2}$ for a positive-measure set of prices, it is easy to check that the distribution over F is different from σ for a positive-measure set of price realizations, a contradiction. Therefore, firms must mix uniformly between $\alpha = 0$ and $\alpha = 1$ in symmetric Nash equilibrium, independently of their realization of p. ■

This result has several noteworthy features. First, equilibrium displays price dispersion in real terms. Both the mean value and the range of real interest rates increase with inflation uncertainty (or, equivalently, with the consumers' coefficient of absolute risk aversion).[4] To see why, suppose that we subject ε to a mean-preserving spread. Because u is concave, c goes up. Note that it is inflation uncertainty (measured by the certainty equivalent c) that affects the structure of equilibrium prices, while expected inflation μ is irrelevant (this follows from the CARA assumption). Second, for any realization of equilibrium real interest rates, consumers are entirely swayed by the frame they adopt. If they think in nominal terms (thus treating an indexed contract as risky), they will select the firm that offers the lower nominal rate, and if they think in real terms (thus treating an indexed contract as safe), they will select the firm that offers the lower real rate. This is because $|p_1 - p_2| < c$ with probability one, such that if the consumer is led to regard one contract as safe and the other contract as risky, he will opt for the former. Third, firms mix uniformly between indexation and no indexation, independently of the real interest rate they adopt.

It may be interesting to interpret Proposition 3 more broadly, in terms of the phenomenon of narrow bracketing of financial risk alluded to at the beginning of this sub-section. The two firms compete in lotteries, and in equilibrium they describe them in a way that either neglects or incorporates the background noise ε, with equal probability and independently of their pricing of the lottery. As the background noise becomes larger, the dispersion of the lottery's expected value goes up, while its expected value for the consumer goes down. In other words, in a competitive market for financial products, greater background risk harms investors who are vulnerable to narrow bracketing.

As in other applications of the competitive framing framework, the equilibrium analysis in this sub-section hinges on the way consumers react to conflicting frames. Here, I assumed that consumers adopt the first frame they encounter. In contrast, if

[4]For existing research on the effects of inflation uncertainty on real price distributions (mark-ups as well as dispersion), see Cukierman (1983) and Benabou and Gertner (1993).

we assumed that consumers adopt a nominal frame whenever it is implied by at least one firm (i.e., when $\alpha_i = 0$ for some firm i), then in symmetric Nash equilibrium, firms would play the contract $(\mu, 0)$. In this case, the market outcome would be competitive in the sense that the expected real interest rate would be zero, but the outcome would not be efficient if we use consumers' expected utility from real outcomes as the welfare criterion, because consumers would be fully exposed to inflation uncertainty. If we assumed that consumers adopt a real frame whenever it is implied by at least one firm, firms would play $(0, 1)$ in symmetric Nash equilibrium, thus coinciding with the benchmark in which consumers are rational and care about real outcomes.

The lessons from this sub-section can be summarized as follows: (i) our framework can capture bracketing of risk as a framing strategy in competitive situations in which firms effectively offer lotteries against some background noise; (ii) the greater the riskiness of the background noise, the higher the expectation and dispersion of equilibrium real prices; (iii) to apply WR, sometimes it is useful to redefine what constitutes alternatives and marketing messages, as the proof of Proposition 3 demonstrates.

4.3 Spurious Product Categorization

In this sub-section I use the formalism to examine the effects of spurious product categorization in a model of price competition with differentiated products. Imagine a consumer choosing between different brands of plain yogurt, which may differ in several objective characteristics such as texture, fat content or sweetness. Suppose further that one producer designs its advertising campaign in a way that positions its brand of yogurt as "dessert", while its rival positions its own brand as a "health" product. This may have several effects. First, the two products are less likely to coexist in the consumer's "consideration set", because when the consumer considers one product, he is more likely to think of alternatives that belong to the same category (see Eliaz and Spiegler (2011)). Second, the different categorization of the two products may accentuate their differences along the objective dimensions, thus making them seem like weaker substitutes than if they were assigned to the same product category. How would these considerations affect the equilibrium pricing of the two brands?

Formally, I begin with a typical "Hotelling" setting. The interval $[1, 2]$ represents a set of possible product types. The two firms sell products that are represented by the two extreme points of this interval - i.e., firm 1 (2) is located at the point 1 (2). The consumer's ideal product type z is distributed according to a continuous and strictly increasing *cdf* $G[1, 2]$, which is symmetric around the interval's midpoint.

Initially, the consumer is randomly assigned to one of the firm (with probability $\frac{1}{2}$ each), independently of his ideal point, and the assigned firm serves as his default option.

Let $A = \mathbb{R}_+$ be the set of feasible product prices, and let $M = \{m, n\}$ be a set of two categories to which each firm can spuriously assign its product. Let $F = \{0, 1\}$, and assume that π assigns probability one to $f = 1$ ($f = 0$) whenever $m_1 = m_2$ ($m_1 \neq m_2$). That is, $f = 1$ (0) means that the two products are assigned to the same category (different categories). For each frame f, let $c_f > 0$ represent the consumer's "transportation cost" - i.e. his rate of substitution between price and product type. Let $\theta_f \leq 1$ represent the probability that the consumer includes firm j's product in his consideration set when he is initially assigned to firm i. The functions c and θ are both symmetric. Assume $c_1 \leq c_0$ and $\theta_1 \geq \theta_0$, with at least one strict inequality. That is, when the two products are identically categorized, the consumer is more likely to consider both of them and he treats them as closer substitutes.

When the consumer's ideal point is $z \in [0, 1]$ and he is initially assigned to firm i, he makes a comparison with probability θ_f. If he does not make a comparison, he chooses firm i's product. Conditional on making a comparison, he chooses firm i if and only if

$$p_i + c_f \cdot |i - z| \leq p_j + c_f \cdot |j - z|$$

Let $p_1 < p_2$. Then,

$$s_1(p_1, p_2, f) = \theta_f \cdot G\left(\frac{3}{2} + \frac{p_2 - p_1}{2c(m_1, m_2)}\right) + (1 - \theta_f) \cdot \frac{1}{2}$$

For simplicity, assume that $G \equiv U[1, 2]$. (The reason I did not impose this restriction at the outset is that the linearity would mask differences between this model and Example 1.3.) It follows that

$$s_1(p_1, p_2, f) = \frac{1}{2}\left[1 + \theta_f \cdot \min(1, \frac{p_2 - p_1}{c_f})\right]$$

We are now ready to characterize the firms' equilibrium pricing and marketing strategies.

Proposition 4 *In Nash equilibrium, firms charge*

$$p^* = \frac{2c_0c_1}{c_0\theta_1 + c_1\theta_0}$$

and randomize uniformly between m and n.

Proof. First, observe that π satisfies WR: if one firm plays m and n with equal probability, the distribution over the consumer's frame is uniform, independently of the other firm's marketing strategy. By Corollary 1, it follows that in Nash equilibrium,

$$s_1^*(p_1, p_2) = \frac{1}{2}\left[1 + \frac{1}{2}\theta_0 \cdot \min(1, \frac{p_2 - p_1}{c_0}) + \frac{1}{2}\theta_1 \cdot \min(1, \frac{p_2 - p_1}{c_1})\right]$$

for almost every price profile (p_1, p_2). Thus, firm 1's market share for almost any (p_1, p_2) for which $p_1 \leq p_2$ is

$$\begin{array}{lll} \frac{1}{2}[1 + (\frac{\theta_1}{2c_1} + \frac{\theta_0}{2c_0})(p_2 - p_1)] & if & p_2 - p_1 \leq c_1 \\ \frac{1}{2}[1 + \frac{\theta_1}{2} + \frac{\theta_2}{2c_2}(p_2 - p_1)] & if & c_1 < p_2 - p_1 \leq c_0 \\ \frac{1}{2}[1 + (\frac{\theta_1}{2} + \frac{\theta_0}{2})] & if & p_2 - p_1 > c_0 \end{array}$$

It is now straightforward to show that firms' pricing decision in Nash equilibrium is as stated in the proposition.

To establish the firms' marketing strategies, suppose that one firm does not randomize uniformly over M - in particular, w.l.o.g, assume that firm 2 plays m with probability greater than $\frac{1}{2}$. By assumption, $\theta_1/c_1 \neq \theta_0/c_0$. Therefore, firm 1 would profit from deviating to the pure strategy $(p^* - \varepsilon, m)$, as long as $\varepsilon > 0$ is sufficiently small. In contrast, when firm 2 randomizes uniformly over M, firm 1's market share is independent of its own marketing strategy, hence uniformly randomizing over M is consistent with best-replying. ■

Let us use this result to perform two simple comparative statics exercises. Suppose that in the original state, $c_1 < c_0$ and $\theta_1 > \theta_0$. First, fix the transportation costs and modify the consideration probabilities into $\theta_0' = \theta_1' = \frac{1}{2}(\theta_0 + \theta_1)$. The interpretation is that we keep the "average" consideration probability constant, while eliminating its dependence on spurious product categorization. It is easy to see that the equilibrium price *rises* as a result. Second, fix the consideration probabilities and modify the transportation costs into $c_0' = c_1' = \frac{1}{2}(c_0 + c_1)$. The interpretation is that we keep "average" substitutability constant, while eliminating its dependence on spurious product categorization. In this case, the equilibrium price *decreases* as a result. In both cases, the equilibrium marketing strategy is not affected by the modification. The lesson is that the effect of spurious categorization on consumer attention lowers the equilibrium price, while its effect on perceived substitutability raises the equilibrium price. This result illustrates the modeling framework's ability to generate insights into various effects of advertising and marketing.

Comment: Models with conventionally rational consumers
A typical reaction to behavioral I.O. models (as well as other models with boundedly rational agents) is whether they are behaviorally equivalent to models with conventionally rational agents. The application in this sub-section is particularly reminiscent of more conventional models of product differentiation and advertising. Therefore, this may be a good point to pause and make a careful comparison.

Product categorization in our model is essentially an advertising strategy that influences the effective substitutability between the two products. Therefore, it is useful to compare the current model to the model of complementary advertising due to Becker and Murphy (1993), according to which consumers would choose between firms as if they maximized a utility function over pairs (p, m). The two models are behaviorally distinct, because the consumer choice function in the current model is inconsistent with utility maximization. To see why, suppose that $c_1 < p_2 - p_1 < c_0$. Then, there exist consumers whose ideal point is sufficiently close to $z = 2$, who would choose firm 1 under $f = 1$ and firm 2 under $f = 0$. This means that their revealed preference relation would be

$$(p_1, m) \succ (p_2, m) \succ (p_1, n) \sim (p_1, m)$$

contradicting rationality.

Let us turn to a comparison between spurious categorization in the current model and conventional product differentiation. Imagine that there are no framing effects, i.e. c is a constant that cannot be manipulated by firms, and $\theta = 1$. Instead, assume that an action in M corresponds to a location decision on the interval $[1, 2]$. In this case, when both firms choose the same location, it is as if $c = 0$. As a result, there exists no pure-strategy Nash equilibrium in this analogous model. In particular, any equilibrium has to admit random pricing. It can be shown that other variations on the spatial-competition analogue would also be behaviorally distinguishable from the model of this sub-section.

4.4 Manipulating Subjective Weights on Product Attributes

When consumers face multi-attribute products or multi-dimensional pricing plans, a natural choice criterion is to maximize some weighted average value across attributes or dimensions. This mode of behavior enters the domain of our present inquiry when firms use marketing devices to manipulate the weights. In this sub-section I examine two examples of models that fall into this category. My objective is to analyze the structure of equilibrium offers and the way they are marketed by firms. In both models, I define $A = [0, \infty)^K$ and $p(a) = 1 - \frac{1}{K}\sum_k a^k$, where a^k is interpreted as the value of product attribute k to the consumer. For every k, let e_k denote the K-vector e for which $e^k = 1$ and $e^j = 0$ for all $j \neq k$.

4.4.1 Framing by Suggesting Weights

Define $M(a) = F = \Delta\{1, ..., K\}$ for all a, and assume that $\pi(m_1, m_2)$ assigns probability one to $f = \frac{1}{2}(m_1 + m_2)$. The choice function is

$$s_1(a_1, a_2, f) = \frac{1}{2}\left[1 + sign\left(\sum_k f^k \cdot (a_1^k - a_2^k)\right)\right]$$

The interpretation is that the consumer's frame consists of the weights he assigns to the different attributes. The firms' marketing messages are suggested weights, and the consumer adopts the average of their suggestions.[5]

In a benchmark case in which consumers apply uniform weights independently of the firms' marketing messages, the model is reduced to simple Bertrand competition, and in Nash equilibrium firms offer alternatives a for which $\sum_k a^k = K$. The following result shows that under the current model of manipulable weights, in Nash equilibrium firms only offer alternatives that assign the maximal value to one attribute and zero value to all other attributes, and their marketing message assigns all weight to this attribute.

Proposition 5 *In any Nash equilibrium, each firm mixes over strategies of the form* $(a, m) = (Ke^k, e^k)$.

Proof. redefine alternatives to consist of only the total value, such that marketing involves both m and the variation of values across attributes. show that a firm that offers a better total value product can enforce being chosen against any inferior alternative, by being maximally attractive on one attribute and putting all weight on it. this means that WIM is satisfied because the strategy max-min market share is one for the better firm. therefore firms choose alternatives as if the weights are uniform. this gives us the rational benchmark in terms of profits. now, assume that one firm

[5] Zhou (2008) analyzes a monopolistic model in which the firm can influence consumers' subjective weights in a similar manner.

plays something other than (Kek,ek) with positive probability. then the other firm, by mixing over all ek-epsilon,ek, gets a positive market share and positive payoff. ■

Thus, competitive forces push firms to offer alternatives with a maximally skewed distribution of values across product attributes, where one attribute is maximally attractive to the consumer and all other attributes are minimall attractive. Firms accompany these offers with marketing messages that try to make the attractive attribute as salient as possible. The equilibrium outcome is competitive in the sense that it generates zero profits, but it is highly obfuscatory in the sense that (for $K > 2$) consumers' subjective evaluation of market alternatives is necessarily higher than their value according to uniform weights. The equilibria in which the difference between subjective and objective values are maximized are those in which firms play the same pure strategy (Ke^k, e^k) with probability one. Note that the proof of this result does not involve WR but the weaker property WIM, because the model does not satisfy WR.

4.4.2 Ahn-Ergin Framing

Let $\Omega = \{1, ..., K\}$ be a set of states of nature, where $K > 1$, and interpret A as a set of feasible acts. The interpretation of the definition of $p(a)$ is that firms hold a uniform prior belief over the state space. Let M be the set of all partitions of Ω, and define $M(a) \subseteq M$ as the set of partitions for which $a^j \neq a^k$ implies that j and k belong to different cells. Let $F = \Delta(\Omega)$. Assume that π assigns probability one to some probability distribution over the state space, which assigns equal weight to all cells in $m_1 \vee m_2$ (the coarsest refinement of m_1 and m_2).

The model of consumer choice is based on Ahn and Ergin (2010), who were in turn motivated by a theory of probability weighting due to Tversky and Koehler (1994) known as "support theory". The interpretation in our context is that firms present acts as functions of mutually exclusive *events*. Thus, when a firm's offer assigns the same value to multiple states, the firm has a degree of freedom in presenting it. For instance, when $K = 3$, the act $(1, 1, 0)$ can be presented as ($a\{1, 2\} = 1$, $a\{3\} = 0$), or as ($a\{1\} = 1$, $a\{2\} = 1$, $a\{3\} = 0$). For a rational consumer, this degree of freedom is irrelevant. In contrast, the psychology underlying the Ahn-Ergin model is that splitting an event into a collection of separate sub-events can increase the consumer's subjective probability of the event, because the detailed enumeration of the sub-events enhances the event's perceived importance. When the firms' partitions overlap, the consumer takes their coarsest refinement as the effective collection of atomic events.

The model is trivial when $K = 2$, because when $a_i^1 \neq a_i^2$ the set of feasible marketing actions for firm i is a singleton (the partition $\{\{1\},\{2\}\}$), while when $a_i^1 = a_i^2$ there is no obfuscation and the consumer's perception of the value of firm i's act is correct, independently of the firms' marketing actions. From now on, assume $K > 2$. To illustrate the consumer's choice function, let $K = 3$. Suppose that firm 1 plays the act $(1, 1, 0)$ and adopts the partition $\{\{1, 2\}, \{3\}\}$. Suppose that firm 2 plays the act $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. Note that the consumer's evaluation of firm 2's act is independent of the partition he employs. However, firm 2's partition will affect his evaluation of firm 1. If firm 2 adopts the same partition as firm 1, the consumer will evaluate firm 1's act at $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 = \frac{1}{2}$, hence he will be indifferent. If firm 2 switches to the partition $\{\{1\}, \{2, 3\}\}$, the join of the firms' partitions is the finest possible partition, in which case firm 1's perceived value will be $\frac{1}{3} \cdot 1 + \frac{1}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3}$, hence the consumer will strictly prefer firm 1.

This model satisfies WR because each firm can unilaterally enforce the uniform prior by playing the finest partition $\{\{1\}, ..., \{K\}\}$, which belongs to $\cap_{a \in A} M(a)$. By Corollary 1, firms' equilibrium choices of acts proceeds as if the consumer has a uniform prior, and therefore $p(a_1) = p(a_2) = 0$ in every Nash equilibrium. In particular, if firms play $a^k = 1$ for all k and adopt the finest partition, this constitutes an equilibrium that involves no obfuscation: the consumer's subjective expected average price is zero, which is also the actual expected price according to the firms' "objective" prior.

The following result characterizes the equilibria that involve the highest amount of obfuscation, in the sense that they maximize the difference between the subjective and objective expected equilibrium price.

Proposition 6 *The Nash equilibria that maximize the consumer's subjective (i.e. perceived) expected value have the following form: there exists a state k, such that both firms play $a^k = Ke^k$ and $m = \{\{k\}, \Omega \backslash \{k\}\}$. Consumers' subjective expected payoff in these equilibria is $\frac{1}{2}K$, whereas their objective expected payoff is 1.*

Proof. Consider an arbitrary prior $(f^1, ..., f^K)$, and w.l.o.g, assume that $f^1 > f^2 \geq \cdots \geq f^K$. Then, the price vector $a = (K, 0, ..., 0)$ maximizes $\sum_k f^k a^k$ subject to the constraint that $\sum_k a^k = K$. Now, fix this vector a and look for the prior f that maximizes $\sum_k f^k a^k$, subject to the constraint that f is induced by some partition $m \in M(a)$. It is easy to see that $f = (\frac{1}{2}, \frac{1}{2(K-1)}, ..., \frac{1}{2(K-1)})$, and the partition that generates f is $m = \{\{1\}, \{2, ..., K\}\}$. By construction, this is a symmetric Nash equilibrium strategy: if a firm deviates it to some other strategy (a', m') with $p(a') > 0$, the consumer uses the partition $m \vee m'$ to construct the prior, and by definition the subjective expected value according to this prior will be below $\sum_k f^k a^k$, hence the consumer will not choose the deviating firm and the deviation will be unprofitable. ■

In this set of equilibria, firms offer the minimal value $a^k = 0$ in every state k except one, where the value is as high as possible to ensure zero profits, and they use the smallest partition that is consistent with it. The consumer assigns equal weight to both cells, such that his subjective expected value is $\frac{1}{2}K$. Note that these equilibria also have the minimal size of the consumer's adopted partition. To see why, suppose that the consumer adopts the degenerate partition $\{\Omega\}$ with positive probability in some Nash equilibrium. Then, each firm must play the following strategy with positive probability: $a^k = 1$ for all k, coupled with and $m = \{\Omega\}$. But this means that if firm 2 deviates to the following pure strategy, $a^1 = K - \varepsilon$, $a^l = 0$ for every $l > 1$ and $m = \{\{1\}, \{2, ..., K\}\}$, it wins the consumer with positive probability one if $\varepsilon > 0$ is sufficiently small, and thus earns positive profits. Thus, in every Nash equilibrium, the consumer's adopted partition has at least two cells with probability one.

Comment on the model's scope

The models in this sub-section share the property that the weights consumers apply to multi-attribute products are exclusively a function of the firm's marketing messages. In the first example, the set of feasible marketing messages is independent of the attribute-specific values the firm offers, whereas in the case of Ahn-Ergin framing, it is not. Kőszegi and Szeidl (2012) construct a model in which consumers' subjective weights are exclusively a function of the quality vector itself, and there is no additional marketing message. Specifically, their model satisfies that $f^k > f^l$ whenever $\left|a_1^k - a_2^k\right| > \left|a_1^l - a_2^l\right|$ - that is, the consumer places a higher weight on attributes that exhibit greater variation across firms. The consumer's decision rule does not involve any distinction between "alternatives" and "marketing messages". If we wanted to accommodate it into our framework, we would have to redefine the primitives in a way that trivializes that distinction, such that $M = A$, $M(a) = \{a\}$ and the consumer's frame is identified with (a_1, a_2). It can be shown that firms earn zero profits in Nash equilibrium when the consumer follows the Kőszegi-Szeidl model.

The lessons from section are as follows: (i) the modeling framework can accommodate situations in which firms offer multi-attribute products and use marketing to influence the decision weights that consumers apply to each attribute; (ii) in the examples we analyzed, the equilibrium market outcome is competitive in the sense that firms earn zero profits, but there can be a large gap between the consumers' subjective evaluation of equilibrium offers and their true value; in the maximally obfuscatory equilibria, firms coordinate on the same vector Ke^k and use marketing to maximize the weight consumers apply to attribute k; (iii) when one firm can always ensure that it will be chosen by the consumer whenever its opponent offers an objectively inferior product, the multi-attribute model satisfies WIM; this implies that firms choose their multi-attribute offers as if consumers choose rationally, and any equilibrium outcome is therefore competitive in the sense that firms earn zero profits.

5 Conclusion

My objective in this paper was to present a framework for modeling market competition that involves utility-relevant aspects (such as price or quality) as well as utility-irrelevant aspects that affect the "frame" of the consumer's choice problem, and show how it can be applied to a wide variety of situations. The concepts of WR and WIM emerged as properties that unify a variety of examples. Hopefully, this variety will convince the reader that the interplay between framing and competition can be fruitfully modeled at a level of generality that abstracts from the concrete psychological mechanisms underlying consumer choice. This abstract approach complements the common practice in behavioral I.O. of focusing on one aspect of consumer psychology at a time.

Although WIM turns out to be useful in a large number of examples, it is not a robust property. For instance, in the model of Section 4.2, suppose that when consumers are presented with both nominal and real frames, they adopt the former with some probability $q \neq 0, \frac{1}{2}, 1$. In this case, WIM ceases to hold, and analysis of equilibrium behavior is an open problem. Looking for alternative properties of frame-sensitive choice that have rich implications for competitive market settings is an important challenge for future work. In particular, we should seek properties with intrinsic behavioral interest. We should also explore properties that are useful when there are $n > 2$ competing firms, or when consumers have an ex-ante outside option; these extensions sever the formal link with zero-sum games, and thus call for new equilibrium characterization techniques.

The model raises several conceptual problems. First, several applications (e.g., Sections 4.1 and 4.2) demanded an explicit assumption regarding consumers' response to conflicting frames. In some cases, it is reasonable to assume that consumers will adopt any of the suggested frames with some probability. In other cases, the multiplicity of frames annuls the framing effect altogether because consumers become "enlightened".

Is there a principled selection of a frame in such situations? New experimental work may illuminate this question.

Second, the model of consumer choice in this paper implicitly assumes that consumers do not know the equilibrium and thus draw no strategic inferences about the value of alternatives from the marketing messages that accompany them. In some cases, this makes sense because the market situation does not confront the consumer with an explicit distinction between alternatives and marketing messages. However, in other cases marketing messages can be viewed as visible and easily identifiable "packages" of utility-relevant, yet difficult-to-evaluate content. In these cases, consumers may use knowledge of the equilibrium correlation between a and m to make better choices. For instance, if symmetric equilibrium in Example 2.2 has the property that expected price conditional on employing the format m is higher than expected price conditional on employing the format n, and if consumers understand this equilibrium correlation, then when they face two realizations $(p, m), (p', n)$, they will use this understanding and choose (p', n), even if m and n are incomparable. Of course, when symmetric equilibrium in our model exhibits no correlation between a and m, it is robust to such inferences. Incorporating these considerations into the model is an interesting direction for future research.

References

[1] Ahn, D. and H. Ergin (2010): "Framing Contingencies," Econometrica 78, 655-695.

[2] Armstrong, M. (2008): "Interactions between Competition and Consumer Policy," Competition Policy International 4, 97-147.

[3] Bachi, B. (2012): "Competition with Price Similarities," mimeo, Tel Aviv University.

[4] Becker, G. and K. Murphy (1993): "A Simple Theory of Advertising as a Good or Bad," Quarterly Journal of Economics 108, 941-964.

[5] Benabou, R. and R. Gertner (1993): "Search with Learning from Prices: Does Increased Inflationary Uncertainty Lead to Higher Markups?" Review of Economic Studies 60, 69-95.

[6] Bernheim, D. and A. Rangel (2009): "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," Quarterly Journal of Economics 124, 51–104.

[7] Carlin, B. (2009): "Strategic Price Complexity in Retail Financial Markets," Journal of Financial Economics 91, 278-287.

[8] Chioveanu, I. (2008): "Advertising, Brand Loyalty and Pricing," Games and Economic Behavior 64, 68-80.

[9] Cukierman, A. (1983): "Relative price variability and inflation: A survey and further results," Carnegie-Rochester Conference Series on Public Policy 19, 103-157.

[10] Eliaz, K. and R. Spiegler (2011): "Consideration Sets and Competitive Marketing," Review of Economic Studies 78, 235-262.

[11] Ellison, G. (2006): "Bounded Rationality in Industrial Organization," in R. Blundell, W. Newey and T. Persson (eds.), Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press.

[12] Eyster, E. and G. Weizsäcker (2011): "Correlation Neglect in Financial Decision-Making," Discussion Paper no. 1104, DIW Berlin.

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

[14] Hart, S. (2008): "Discrete Colonel Blotto and General Lotto Games," International Journal of Game Theory 36, 441–460.

[15] Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," Econometrica 47, 263-291.

[16] Kahneman, D. and A. Tversky (2000): Choices, Values, and Frames, Cambridge University Press, Cambridge.

[17] Kőszegi, B. and A. Szeidl (2012): "A Model of Focusing in Economic Choice," mimeo, UC-Berkeley.

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

[19] Ok, Ortoleva and Riella (2011a): "Revealed (P)Reference Theory," mimeo, NYU.

[20] Ok, Ortoleva and Riella (2011b): "Theory of Product Differentiation in the Presence of the Attraction Effect," mimeo, NYU.

[21] Piccione, M. and R. Spiegler (2012): "Price Competition under Limited Comparability," Quarterly Journal of Economics 127, 97-135.

[22] Rabin, M. and G. Weizsäcker (2009): "Narrow Bracketing and Dominated Choices," American Economic Review 99, 1508-1543.

[23] Salant, Y. and A. Rubinstein (2008): "(A,f): Choices with Frames," Review of Economic Studies.75, 1287-1296.

[24] Shafir, E., P. Diamond and A. Tversky (1997): "Money Illusion," Quarterly Journal of Economics 112, 341-374.

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

[26] Spiegler R. (2011): Bounded Rationality and Industrial Organization. Oxford University Press, New York.

[27] Tversky, A. and D. Koehler (1994): "Support Theory: A Nonextensional Representation of Subjective Probability," Psychological Review 101, 547–567.

[28] Zhou, J. (2008): "Advertising, Misperceived Preferences, and Product Design," mimeo, UCL.

Appendix: Completed Proof of Proposition 3

Consider the symmetric two-person game in which players' action space is $[0, \infty)$ and player i's payoff is $p_i \cdot s_i^*(p_1, p_2)$, where $s_1^*(p_1, p_2)$ is given by (1), and $s_1^* + s_2^* \equiv 1$.

I am grateful to Beniamin Bachi for suggesting the equilibrium strategy below. Moreover, the proof of the result below is a direct extension of the proof technique in Bachi (2012), who analyzed price competition when consumers are unable to distinguish between "similar" prices, and contains the definition of prices as similar if the difference between them is less than c as a special case. The market share function given by (1) can be re-interpreted in terms of Bachi's model, as a case in which half the consumer population have $c = 0$ and the other half have some $c > 0$.

Proposition 7 *The game has a unique symmetric Nash equilibrium, in which each player plays the mixed strategy given by the cdf* $G(p) = \frac{3}{2}(1 - \frac{c}{2p})$ *over the interval* $[\frac{c}{2}, \frac{3c}{2}]$.

Proof. First, note that the equilibrium strategy cannot contain atoms, by standard undercutting arguments. Denote the *cdf* that represents the equilibrium mixed strategy by G. Let p_l and p_h denote the minimal and maximal prices in the support of G. By definition, p_l and p_h are best-replies against G, and in particular weakly more profitable than the prices $p_l + c$ and $p_h - c$. Therefore:

$$\begin{aligned} p_l \cdot [1 - \frac{1}{4}G(p_l + c)] &\geq (p_l + c) \cdot [\frac{1}{2}(1 - G(p_l + c)) + \frac{1}{4}] \\ p_h \cdot [\frac{1}{4}(1 - G(p_h - c))] &\geq (p_h - c) \cdot [\frac{1}{2}(1 - G(p_h - c)) + \frac{1}{4}(1 - G(p_h - c))] \end{aligned}$$

Note that the expressions on the R.H.S of these two inequalities are lower bounds on the profits generated by the prices $p_l + c$ and $p_h - c$. By simple algebra, it follows that $p_l \geq \frac{1}{2}c$ and $p_h \leq \frac{3}{2}c$. Therefore, in particular, $p_h - p_l \leq c$. However, if $p_h - p_l < c$, then a firm can profitably deviate to $p_h + \varepsilon$, if $\varepsilon > 0$ is sufficiently small, because its market share would be $\frac{3}{4}$ both before and after the deviation. It follows that $p_l = \frac{1}{2}c$ and $p_h = \frac{3}{2}c$. Therefore, the market share generated by p_h is precisely $\frac{1}{4}$, such that the equilibrium payoff is pinned down by $\frac{3}{8}c$. The payoff from any p in the support of G thus satisfies

$$\frac{3c}{8} = p \cdot [\frac{1}{2}(1 - G(p)) + \frac{1}{4}]$$

and this pins down the expression for G. If the support of G is not connected, then G must have an atom, a contradiction, hence the support of G is $[\frac{c}{2}, \frac{3c}{2}]$. Checking that deviations to prices outside the support are unprofitable is straightforward. ■