Competitive Framing[*]

Ran Spiegler[†]

March 30, 2013

Abstract

I present a simple framework for modeling two-firm market competition when consumer choice is "frame-dependent", and firms use costless "marketing messages" to influence the consumer's frame. This framework embeds several recent models in the "behavioral industrial organization" literature. 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 variety of applications.

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.

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 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 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 a manipulable parameter in the consumer's probabilistic choice function.

The modeling framework can capture a variety of market situations in which firms manipulate consumer choices - e.g., 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 of framing in competitive marketing settings, such as bracketing of financial risk or spurious product categorization that influences consumer attention and 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 the firms' prices p_1, p_2 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 $p_i < p_j$, firm i's market share under the frame f is $\frac{1}{2}\{1 + G[(p_j - p_i)/(1 - f)]\}$.

How can firms manipulate the consumer's frame? The set of feasible marketing messages given p is $\{Lg \leq p \mid L \in \mathbb{Z}\}$, where $g > 0$ is exogenously given and arbitrarily small. 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, and g represents the density of the grid on which the origin can be located. 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 the consumer's choice function may satisfy. We will say that π satisfies "*Weighted*

3

Regularity" (**WR**) if there is a distribution σ over F, such that for every alternative a that a firm may offer it can find a mixture λ_a over $M(a)$ that induces the distribution σ over F, independently of the other firm's marketing message. WR means that each firm can unilaterally impose the distribution σ over F, thus rendering consumer choice neutral to the rival firm's marketing strategy. 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.

The basic result of this paper is that if WR holds, then firms' market shares in Nash equilibrium are almost surely as if the distribution over F is exogenously set to σ. 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 almost every realization (a_1, a_2) are determined by the value of the associated zero-sum game.

Although WR is not a "generic" property, it turns out to hold in surprisingly many applications. Observing that WR holds greatly facilitates equilibrium analysis, and often leads to 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, a firm's optimal price against the opponent's equilibrium strategy can be computed as if $f = 0$ with probability one. 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 away from the default frame. In this sense, competition disciplines framing (by comparison, a monopolist facing a consumer with some exogenous outside option would strictly prefer to manipulate the 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 given the constraints on consumer rationality.

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 can effectively screen the consumer's type and approximate the first-best, using a random strategy that involves "complex" pricing.

So far, Behavioral Industrial Organization 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 specific psychological phenomena. 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 and the concept of WR to a much wider class of frame-sensitive choice procedures, we can analyze market implications of many phenomena of consumer psychology.

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 the choice problem's frame (more concretely, whether one of the feasible alternatives is a status quo). Salant and Rubinstein (2008) and Bernheim and Rangel (2009) 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 $\pi : M \times M \to \Delta(F)$, 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 $p : A \to \mathbb{R}$ that specifies the net profit to the firm from selling 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 symmetric market share function that satisfies the no-outside-option assumption. To see why, note that we could always set $M(a) \equiv \{a\}$ and let F be a singleton. 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 framework lies in the language it provides for unifying Behavioral Industrial Organization models, as well as in the fruitful restrictions on choice behavior that it suggests.

The following are examples of models that fit conveniently 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 $p(a^1, a^2) = 1 - (a^1 + a^2)$. Let $F = M(\cdot) = \{0, 1\}$. The interpretation of $m = 1$ is that the firm "shrouds" the second attribute, and $f = 1$ means that the second attribute ends up being shrouded. Let $\pi_1(m_1, m_2) \equiv m_1 m_2$. 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, 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' formats are comparable. 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]

---

[1] Carlin (2009) and Chioveanu and Zhou (2011) analyze special cases of π and extend the choice model to $n > 2$ firms. Kamenica et al. (2011) analyze an asymmetric game in which each firm is restricted to a different price format - fixed monthly payment vs. price per minute (consumers' choice function in this model is sensitive to cardinal effective price differences).

Example 2.3: Using partitions to frame acts (based on Ahn and Ergin (2010))

Let $\Omega = \{1, ..., K\}$ be a set of states of nature, and let $A = [0, \infty)^K$ be a set of feasible acts (with monetary consequences). Define $p(a) = 1 - \frac{1}{K}\sum_k a^k$; the interpretation is that firms are risk-neutral and have a uniform prior over Ω. 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 $F = M$ and assume that π assigns probability one to $m_1 \vee m_2$ (the coarsest refinement of m_1 and m_2). Define $\beta(f) \in \Delta(\Omega)$ as the distribution over Ω that assigns equal weight to all cells in f, and equal weight to all states within any given cell. The choice function selects the alternative a that maximizes $\sum_k \beta^k(f)a^k$.

The model of consumer choice is adapted from 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 enhances the event's perceived importance and therefore increases the consumer's subjective probability of it. When the firms' partitions overlap, the consumer takes their coarsest refinement as the effective collection of atomic events.

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 a_i corresponds to player i's type, such that a state consists of the pair (a_1, a_2); player i's (type-dependent) action set is $M(a_i)$; player i's payoff function is s_i^* (m_1 and m_2 are the players' actions, while a_1 and a_2 serve as parameters); and 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.

Comment: 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, or a scale as in Example 1.1. In other cases (arguably Example 2.2), 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 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. For instance, in Example 2.1, splitting a product into two attributes may be part of its marketing. 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. I will occasionally draw the demarcating line according to considerations of analytic convenience, overriding a more natural distinction.

3 Weighted Regularity

This section identifies a property that the function π may satisfy, and shows its role in facilitating Nash equilibrium analysis.

Definition 1 *The function π satisfies **Weighted Regularity (WR)** if there exist $\sigma \in \Delta(F)$ and a collection $(\lambda_a)_{a \in A}$, $\lambda_a \in \Delta(M(a))$, such that for every $a \in A$,*

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

for every $n \in M$.

WR means that each firm can unilaterally enforce a given distribution over the consumer's adopted frame, thus rendering the opponent indifferent to marketing. This concept 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.

Examples 1.1, 2.1 and 2.3 trivially satisfy WR, because each firm can unilaterally enforce a deterministic frame ($f = 0$ in the first two examples, and $f = \{\{1\}, ..., \{K\}\}$ in the third example). On the other hand, in Example 2.2, if M contains two messages m, n such that $\pi_1(m, m') > \pi_1(n, m')$ for every $m' \in M$, then π violates WR.

We are now able to state a basic result that will serve us in the applications.

Proposition 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).

Thus, under WR, firms' equilibrium market shares are as if the distribution over the consumer's frame is exogenously given by σ, except possibly for a zero-probability set of profiles of alternatives. The proof of this result will be a corollary of a more general proposition presented in Section 5 (all proofs are relegated to the Appendix). Note that Proposition 1 does *not* imply that the entire game can reduced to a simpler game in which each firm chooses a_i to maximize $p(a_i)s_i^*(a_1, a_2)$. The reason is that the result pins down market shares for profiles (a_1, a_2) that belong to the support of the equilibrium distribution, but not outside it.

4 Applications

In this section I analyze several market models with boundedly rational consumers. In all cases, observing that WR holds (sometimes under a suitable re-specification of the model's primitives) simplifies the equilibrium analysis and leads to a novel economic insight.

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 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 description modes at the same time, most of them would probably have seen through the trick and switched to a more consequentialist way 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. In this sub-section I ask whether competitive forces rule out the elicitation of framing effects, under the assumption that the presence of conflicting frames annuls the framing effect.

Let A be some uncountably infinite set. Let $M(a) = F$ for all $a \in A$. One of the elements in F, denoted 0, is referred to as the "default frame". Assume that $\pi_m(m, m) = 1$ for every m, and $\pi_0(m, n) = 1$ whenever $m \neq n$. Finally, assume that whenever $a_1 \neq a_2$, $f \neq g$ implies $s(a_1, a_2, f) \neq s(a_1, a_2, g)$.

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

Thus, symmetric Nash equilibrium in this game (if one exists) 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 (without transaction costs) belong to this category. In the current model, each firm can unilaterally enforce the default frame, just as in models of speculative trade, each trader can unilaterally prevent trade. In both models, equilibrium selects this unilaterally enforceable action.

Of course, this result is limited because it is stated for a certain class of equilibria that need not exist. Let us now revisit Example 1.1, which described 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.

Proposition 3 *In Example 1.1, there is a unique Nash equilibrium, in which firms play $p = k$ and $m = 0$.*

This example conveys a subtle lesson regarding the interplay between competition and framing effects. The zero-sum aspect of the competitive interaction, and the assumption that each firm can unilaterally enforce the default frame, imply that firms refrain from manipulating the consumer's frame in equilibrium. In this sense, competition does rule out the elicitation of framing effects. By comparison, 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 disciplines framing does not imply that the outcome is more competitive than if consumer choices were manipulated.

There are examples that fall into the default-frame model, in which firms do not necessarily play $m = 0$ in all equilibria. In Example 2.1, the pure strategy $(a, m) = ((-1, 1), 1)$ is an equilibrium strategy. In particular, firms do not have an incentive to unshroud the second attribute because the alternatives are as competitive as possible, whether the consumer considers both attributes or only the first attribute. Thus, although firms' equilibrium payoffs are as if the second attribute is unshrouded, it does not necessarily follow that in equilibrium firms refrain from shrouding.

4.2 Bracketing Financial Risk

In this sub-section I study a model in which firms compete in lotteries and the carrier of the (risk averse) consumer's vNM utility function can be manipulated through framing. It has been observed in experiments (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 (a phenomenon known as "narrow bracketing"), 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 in a competitive market setting. I use the interest-cum-inflation story, but I treat "inflation" additively rather than multiplicatively, for simplicity (this does not affect the qualitative results). 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's 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. In contrast, no indexation ($\alpha = 0$) implies a "nominal frame", thus encouraging the consumer to ignore inflation and think about outcomes in nominal terms. 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 - i.e., real or nominal wealth - is frame-dependent. 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.

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). Thus, because framing manipulates the domain to which the consumer applies his risk preferences, it can lead to preference reversals.

Proposition 4 *The game has a unique symmetric Nash equilibrium. Firms randomize over the expected real interest rate p according to the cdf*

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

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

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).[2] 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 CARA). 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. Thus, equilibrium exhibits spurious multiplicity of contractual forms.

The broad economic lesson from this application is that in a competitive market for financial products, greater background risk harms investors because they are more

---

[2]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).

vulnerable to narrow bracketing. Real price dispersion and multiplicity of contractual forms are observable manifestations of this type of framing. At the methodological level, one lesson is that in order to apply WR, sometimes it is useful to redefine alternatives and marketing messages. In the proof of Proposition 4, I redefine an alternative as the expected real interest the loan bears, and the marketing message is identified with the indexation decision. The consumer's frame takes the values $0, -c, c$, i.e. it is identified with the relative perceived riskiness of the two loans. WR becomes applicable under this re-specification, and this greatly simplifies the proof.

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

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 less likely to think about another product if it belongs to a different 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 motives affect the equilibrium pricing of the two brands?

I use a typical "Hotelling" setting. The two firms sell products that are represented

by the two extreme points of the interval $[1, 2]$ - 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 $\pi_1(m_1, m_2) = 1$ (0) if $m_1 = m_2$ ($m_1 \neq m_2$). The interpretation of $f = 1$ (0) is 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 $i \neq j$. 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,

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

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 1.1.) It follows that

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

Proposition 5 *The game has a unique Nash equilibrium: firms charge*

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

and randomize uniformly between m and n.

Let us use this result to perform two simple comparative statics exercises. Suppose that in the original state, $c_1 < c_0$ and $\theta_1 > \theta_0$. 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. Alternatively, fix the original 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.

Comment: Models with conventionally rational consumers

A typical question concerning 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 advertising. For instance, consider the model of complementary advertising due to Becker and Murphy (1993), according to which consumers choose between firms as if they maximize 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.

5   Worst-Case Independence Of Marketing

In this section I discuss a weaker property than WR, which is defined in terms of the market share function s^*.

Definition 2 *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 $a_2 \in A$.

WIM means that a firm's worst-case-optimal marketing strategy is independent of the alternative that its opponent offers. It is a weaker property than WR, as the following lemma establishes.

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

The following result establishes that under WIM, equilibrium market shares for almost any (a_1, a_2) are given by $v(a_1, a_2)$, the value of the zero-sum game associated with (a_1, a_2). 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 follows as a corollary of this result.

Proposition 6 *Suppose that s^* satisfies WIM. Then, in Nash equilibrium, firm 1's expected market share is $v(a_1, a_2)$ conditional on almost every (a_1, a_2).*

The applications in Section 4 consistently applied Proposition 1. While WIM is a less elegant property than WR, it has a few attractive features. First, WIM is stated in terms of observable variables, namely $(a_i, m_i)_{i=1,2}$, whereas WR is stated in terms of the consumer's frame, which is unobservable in many applications. Moreover, as we saw in Section 4.2, a given model can be formulated in different ways that are equivalent in terms of the market share function and the firms' payoff functions, yet they are based on different specifications of F and different ways to distinguish alternatives from marketing messages. As a result, whether WR holds may depend on the exact specification of F, whereas WIM is not sensitive to it. This makes WIM useful as a diagnostic tool, as the following example illustrates.

Example 5.1: 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 firm's profit conditional on being chosen is the mean of its distribution. Given a price distribution, identify the alternative a with the mean of the distribution, and the marketing message m with the distribution of deviations from the mean. There are various ways to define F and π that would give rise to the same market share function. This raises the question whether one of these specifications might satisfy WR. The answer turns out to be negative, as the market share function violates WIM. Observe that the ex-post zero-sum game associated with any (a_1, a_2) is what Hart (2008) called 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. Since WR implies WIM, it follows that no specification of F and π would satisfy WR.

Finally, there are cases where WIM clearly holds and Proposition 6 can be readily applied, even if WR does not hold. The following lemma addresses a special case of interest. it shows that when an individual firm can ensure beating its rival whenever it offers a more profitable alternative (and gets a market share of at least 50% when the rival offers an equally profitable alternative), any equilibrium outcome must be competitive in the sense that firms only offer alternatives that generate zero profits.

Lemma 2 *Suppose that for every $p \in \mathbb{R}$ there exists (a, m) satisfying $a \in A$, $p(a) = p$ and $m \in M(a)$, such that: (i) $s_1^*((a, m), (a', m')) = 1$ for every (a', m') with $p(a') > p$ and $m' \in M(a')$; (ii) $s_1^*((a, m), (a', m')) \geq \frac{1}{2}$ for every (a', m') with $p(a') = p$ and $m' \in M(a')$. Then, in Nash equilibrium, each firm assigns probability one to alternatives a for which $p(a) = 0$.*

Using this lemma, it immediately follows that firms offer zero-profit alternatives in Examples 2.1 and 2.3. The reason is that in both cases, an individual firm can enforce a frame that induces "rational choice" (in the sense that the consumer necessarily chooses the alternative a with the lower $p(a)$), hence the antecedent of Lemma 2 holds. The following is a richer application of the lemma.

Example 5.2: Manipulating subjective weights on product attributes

Let $A = [0, \infty)^K$ represent a set of multi-attribute products, where $K > 2$, and a^k is interpreted as the value of product attribute k to the consumer. Let $p(a) = 1 - \frac{1}{K} \sum_k a^k$. 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. Let e_k denote the K-vector e for which $e^k = 1$ and $e^j = 0$ for all $j \neq k$.[3]

This model violates WR, and it is not obvious how to recover this property by some re-specification of the primitives. Things are much easier with WIM. Redefine an alternative as $a^* = 1 - p(a)$, and redefine the marketing message as $(m, (a^k - a^*)_{k=1,...,K})$. That is, an alternative is now defined as the average quality the firm delivers to the consumer, and the marketing message consists of the suggested attribute weights as well as the variation of quality across product attributes. It is easy to verify that whenever $a_1^* > a_2^*$, firm 1 can ensure being chosen with probability one, by playing $m = e^k$ and $a^j = 0$ for all $j \neq k$. This leads to the following equilibrium characterization.

Proposition 7 *In any Nash equilibrium, each firm mixes over strategies of the form* $(a, m) = (e^k, e^k)$.

Thus, competitive forces push firms to offer alternatives with a maximally skewed distribution of values across product attributes, with zero quality in all attributes except one. 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 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 is maximized are those in which firms play the same pure strategy (e^k, e^k) with probability one.[4]

---

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

[4]Kőszegi and Szeidl (2012) construct another model of choice between multi-attribute products, in which consumers' subjective weights are exclusively a function of the quality vector itself. There is no

6 Conclusion

My objective in this paper was to present a framework for modeling market competition when firms' competitive strategy 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. The concept of WR emerged as property that unifies a variety of market situations and facilitates their analysis. 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 approach complements the common practice in Behavioral Industrial Organization of focusing on one aspect of consumer psychology at a time.

Although WR turned 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, WR ceases to hold, and analysis of equilibrium behavior is an open problem. One interesting direction for future research is to investigate when WR emerges as a necessary consequence of a larger model that endogenizes the consumer's sensitivity to framing.

The modeling framework raises other challenges for future research. First, a number of applications (e.g., Sections 4.1 and 4.2) demanded an explicit assumption regarding consumers' response to conflicting suggested frames. In some cases, it is reasonable to assume that consumers will adopt any of the suggested frames with some probability. In other cases, multiplicity of frames annuls the framing effect altogether because consumers become "enlightened". Is there a principled criterion for selecting the 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 draw no strategic inferences about the value of alternatives from the marketing messages that accompany them. In some cases, this makes sense because there the market does not provide an explicit distinction between alternatives and marketing messages. However, in other cases marketing messages can be viewed as visible and easily identifiable "packages" of difficult-to-evaluate content. In these cases, consumers may use knowledge of the equilibrium correlation between a and m to make better choices. When an equilibrium exhibits no correlation between a and m (as in Section 4.2), it is robust to such inferences. It would be interesting to incorporate these considerations into the modeling framework.

---

distinct marketing message. Lemma 2 is applicable to their model implies zero equillibrium profits.

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: Proofs

Lemma 1

Consider the zero-sum game induced by a given profile (a_1, a_2). Suppose that firm i mixes over $M(a_i)$ according to the strategy λ_{a_i} as in Definition 1. Then, since this enforces a distribution over f that is independent of firm j's marketing message, firm j is indifferent among all marketing messages, hence $(\lambda_{a_i}, \lambda_{a_j})$ is a Nash equilibrium in the zero-sum game for any a_j. By the Minimax Theorem, this means that λ_{a_i} max-minimizes firm i's expected market share for all a_j, hence Strong WIM holds. ∎

Proposition 6

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). ∎

Proposition 1

Since π satisfies WR, Lemma 1 implies that s^* satisfies Strong WIM. By Proposition 6, in Nash equilibrium, market shares for almost every (a_1, a_2) are given by the value of the associated zero-sum game. As we observed in the proof of Lemma 1, λ_{a_1} as defined in Definition 1 max-minimizes firm 1's market share for every a_2. By definition, λ_{a_1} induces the distribution σ over F, independently of firm 2's marketing message. It follows that

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

for almost every (a_1, a_2). ∎

Proposition 2

Each firm can unilaterally induce the frame $f = 0$, by playing $m = 0$. Therefore, WR trivially holds. By Proposition 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. ∎

Proposition 3

Fix a Nash equilibrium. As already observed, this example satisfies WR. Hence, by Proposition 1, firm i's equilibrium market share for almost any (p_1, p_2) is as if $f = 0$. Therefore, every p_i in the support of each firm i's marginal equilibrium distribution over prices maximizes $p_i \cdot \frac{1}{2}[1 + (Ep_j - p_i)/k]$, where the expectation operator is w.r.t firm j's marginal equilibrium distribution over prices. It follows that $p_i = \frac{1}{2}[k + Ep_j]$ for every i, $j \neq i$, hence both firms choose $p = k$ with probability one in Nash equilibrium. 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 $\lambda \in \Delta\{Lg \leq k \mid L \in \mathbb{Z}\}$ that assigns positive probability to some $m \neq 0$. Suppose that firm 1 deviates to the pure strategy $(k - \varepsilon \cdot sign(m), m)$, where $\varepsilon > 0$. Note that whenever firm 2's realized marketing message is equal to (different from) m, the induced frame is m (0). Moreover, the event in which both firms play m occurs with positive probability as a

result of firm 1's deviation. It is now straightforward to calculate that if ε is sufficiently small, firm 1's deviation is profitable. ∎

Proposition 4

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 $p = r - (1 - \alpha)\mu$, 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.

Suppose that the marginal price distribution p induced by some symmetric equilibrium strategy contains an atom on some p. Then, an individual firm can profitably deviate by playing the price $p - \varepsilon$, where $\varepsilon > 0$ is arbitrarily small, coupled with the same distribution over α that accompanied p prior to the deviation. It follows that the marginal equilibrium price distribution is non-atomic.

Note that π satisfies WR: this property is verified by $\lambda_p(0) = \lambda_p(1) = \frac{1}{2}$ for every p, and $\sigma = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. Therefore, by Proposition 1, in Nash equilibrium the market share function is

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

for almost every price profile (p_1, p_2) in the support of the equilibrium marginal price distribution for which $p_1 < p_2$ (s_2^* is defined symmetrically). Suppose that firms play $\alpha = 1$ with a probability different than $\frac{1}{2}$ for a positive-measure set of prices. Specifically, consider two prices p, p' in the support of the equilibrium marginal price distribution, for which $0 < p' - p < c$. Suppose that the probability $\alpha = 1$ is played conditional on p (p') is $q \neq \frac{1}{2}$ ($q \neq \frac{1}{2}$). Then, the probability of $f = 0$ conditional on the profile (p, p') is $qq' + (1 - q)(1 - q') \neq \frac{1}{2}$, hence $s_1^*(p, p') \neq \frac{3}{4}$. We have established that the probability of such price pairs p, p' to be zero. Therefore, firms must mix uniformly

between $\alpha = 0$ and $\alpha = 1$ for almost every p in the support of the equilibrium strategy. This means that the market share function is given by (1) for every p_1 (including values that lie outside the support of the equilibrium strategy) and almost every p_2. It follows that in order to find the marginal equilibrium *cdf* G over prices, we can regard the model as a price competition game in which each firm i chooses p_i to maximize $p_i \cdot s_i^*(p_i, p_j)$, where s_1^* is given by (1). Let us derive G.[5]

We have seen that G does not contain any atom. 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, $p_h - p_l \leq c$. If this inequality is strict, 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 $\frac{3}{8}c$. 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. ∎

Proposition 5

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

---

[5]I am grateful to Beniamin Bachi for suggesting the equilibrium strategy G in the reduced price competition game. Moreover, my derivation of G is a direct extension of the proof technique in Bachi (2013), who analyzes 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 $c > 0$.

marketing strategy. By Proposition 1, it follows that in Nash equilibrium,

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

for almost every price profile (p_1, p_2) in the support of the equilibrium distribution.

We need to distinguish between two cases. First, suppose that $p_1 \neq p_2$ with positive probability. Then, in order for s_1^* to be given by (2) for some (p_1, p_2) with $p_1 \neq p_2$, it must be the case that each firm i randomizes uniformly over M conditional on charging p_i (unless $|p_2 - p_1| \geq c_0$ and $\theta_0 = \theta_1$, in which case the market share is independent of the firms' marketing messages). It follows that we can assume that firms randomize uniformly over M with probability one. Therefore, for *every* p_i (including values that lie outside the support of firm i's equilibrium strategy), s_i^* is given as follows:

$$s_i^*(p_1, p_2) = \begin{cases} \frac{1}{2}[1 + (\frac{\theta_1}{2c_1} + \frac{\theta_0}{2c_0})(p_j - p_i)] & if & p_j - p_i \leq c_1 \\ \frac{1}{2}[1 + \frac{\theta_1}{2} + \frac{\theta_2}{2c_2}(p_j - p_i)] & if & c_1 < p_j - p_i \leq c_0 \\ \frac{1}{2}[1 + (\frac{\theta_1}{2} + \frac{\theta_0}{2})] & if & p_j - p_i > c_0 \end{cases}$$

It is now straightforward to show that firms' pricing decision in Nash equilibrium is as stated in the proposition, hence $p_1 = p_2$ with probability one, a contradiction.

Now suppose that $p_1 = p_2 = p$ with probability one in equilibrium. Therefore, expression (2) does not convey any information regarding the firms' marketing strategies. Suppose w.l.o.g that firm i plays m with probability greater than $\frac{1}{2}$. By assumption, $\theta_1/c_1 \neq \theta_0/c_0$. Therefore, firm j would profit from deviating to either of the pure strategies $(p - \varepsilon, m)$ or $(p + \varepsilon, n)$, $n \neq m$, as long as $\varepsilon > 0$ is sufficiently small. In contrast, when firm i randomizes uniformly over M, firm j's market share is independent of its own marketing strategy, hence uniformly randomizing over M is consistent with best-replying. Having pinned down the firms' marketing strategy, the only value of p that is consistent with mutual best-replying is p^* as stated in the proposition. ∎

Lemma 2

Let us first verify that there exists an equilibrium in which firms play $p(a) = 0$ with probability one. By assumption, there exists (a, m) with $p(a) = 0$ such that $s_1^*((a, m), (a', m')) = 1$ whenever $p(a') > p(a)$. If both firms play this strategy, there is no profitable deviation for any firm.

Let us now show there exist no equilibria in which firms offer a with $p(a) > 0$ with positive probability. Redefine alternatives and marketing messages as follows: $\tilde{a} = p(a)$ and $\tilde{m} = (a, m)$. By assumption, if a firm offers $\tilde{a}$, it can accompany this choice with

some $\tilde{m}$ such that its market share is one whenever the opponent plays $\tilde{a}' > \tilde{a}$, and at least $\frac{1}{2}$ whenever the opponent plays $\tilde{a}$ as well. This implies that $\tilde{m}$ max-minimizes the firm's market share against any $\tilde{a}'$, hence WIM holds. By Proposition 6, market shares in any Nash equilibrium are as if the consumer chooses the firm that offers the lower $\tilde{a}$ (with a symmetric tie-breaking rule) for almost every $(\tilde{a}_1, \tilde{a}_2)$ in the support of the equilibrium distribution.

Consider the highest value of $\tilde{a}_i$ in the support of firm i's equilibrium strategy, and suppose $\tilde{a}_i > 0$. W.l.o.g, assume $\tilde{a}_1 \geq \tilde{a}_2$. If firm 2's strategy assigns an atom to $\tilde{a}_1$, then firm 1 can deviate by undercutting to $\tilde{a}_1 - \varepsilon$, where $\varepsilon > 0$ is arbitrarily small, coupled with a suitable marketing strategy $\tilde{m}'$ that would ensure that the firm is chosen if 2 plays any $\tilde{a} > \tilde{a}_1 - \varepsilon$ (by assumption, such a marketing strategy exists). We already observed that prior to 1's deviation, 1's market share is zero for almost every profile $(\tilde{a}_1, \tilde{a}_2)$ for which $\tilde{a}_2 < \tilde{a}_1$. Therefore, firm 1's deviation is profitable. If, however, firm 2's strategy does not assign an atom to $\tilde{a}_1$, firm 1's market share when it plays $\tilde{a}$ is zero. Therefore, firm 1's equilibrium profit is zero. If $\tilde{a}_2 > 0$, firm 1 can profitably deviate by playing some $\tilde{a} \in (0, \tilde{a}_2)$, coupled with a suitable marketing strategy $\tilde{m}'$ that would ensure that the firm beats firm 2 whenever it plays $\tilde{a}'' > \tilde{a}$. It follows that $\tilde{a}_2 = 0$. If $\tilde{a}_1 > 0$, firm 2 can profitably deviate in the same manner. Therefore, $\tilde{a}_1 = \tilde{a}_2 = 0$. ∎

Proposition 7

As pointed out in the text, this model satisfies the antecedent of Lemma 2. Therefore, in any Nash equilibrium, firms assign probability one to alternatives a for which $\sum_k a^k = 1$, and earn zero profits. Suppose that firm 1, say, assigns positive probability to some (a, m) satisfying $a \notin \{e^1, ..., e^K\}$. Consider a deviation for firm 2 to a mixed strategy that randomizes uniformly over all strategies $(e^k - (\varepsilon, 0, ..., 0), e^k)$, where $\varepsilon > 0$ is arbitrarily small. Because firm 2's strategy is symmetric across attributes, we can enumerate attributes w.l.o.g such that $1 > a^K \geq \cdots \geq a^1 \geq 0$. The strategy $(e^1 - \varepsilon, e^1)$ for firm 2 beats (a, m) if and only if

$$\frac{1 + m^1}{2}(1 - \varepsilon - a^1) - \sum_{k>1} \frac{m^k}{2} a^k > 0$$

The L.H.S of this inequality is weakly greater than $\frac{1}{2}\left[(1 - \varepsilon - a^1) - a^K\right]$. Recall that $K > 2$ and $a \neq e^K$. By the above definition of a^1 and a^K, $a^1 + a^K < 1$. Therefore, if ε is sufficiently small, the last expression is strictly positive. It follows that firm 2's deviation is profitable. ∎