Price Competition under Limited Comparability*

Michele Piccione† and Ran Spiegler‡

First version: May 2009
This version: June 2011

Abstract

This paper studies market competition when firms can influence consumers' ability to compare market alternatives, through their choice of price "formats". In our model, the ability of a consumer to make a comparison depends on the firms' format choices. Our main results concern the interaction between firms' equilibrium price and format decisions and its implications for industry profits and consumer switching rates. In particular, market forces drive down the firms' profits to a "constrained competitive" benchmark if and only if the comparability structure satisfies a property which we interpret as a form of "frame neutrality". The same property is necessary for equilibrium behavior to display statistical independence between price and format decisions. We also show that narrow regulatory interventions that aim to facilitate comparisons may have an anti-competitive effect.

*A former version of this paper, henceforth referred to as Piccione and Spiegler (2009), was circulated under the title "Framing Competition". We thank Noga Alon, Eddie Dekel, Kfir Eliaz, Sergiu Hart, Emir Kamenica, Ariel Rubinstein, Jakub Steiner, Jonathan Weinstein, numerous seminar participants, and especially the editor and the referees of this journal. Spiegler acknowledges financial support from the European Research Council, Grant no. 230251, as well as the ESRC (UK).

†London School of Economics. E-mail: m.piccione@lse.ac.uk.

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

1 Introduction

Standard models of market competition assume that consumers rank all the alternatives they are aware of. Their preferences may reflect informational constraints, but are complete nonetheless. In reality, consumers are often unable to compare alternatives. Moreover, whether consumers are able to make comparisons often depends on how alternatives are described, or "framed":

- Prices and quantities may be stated in units of measurement that consumers find difficult to compare. For example, interest on a bank deposit can be presented for different time units, nutritional contents of a food product can be specified for various units of weight or volume, etc.

- Price schedules in certain industries contain a large number of contingencies. For instance, a fee structure for banking services specifies different fees for different classes of transactions. Similarly, a calling plan conditions rates on the destination, according to some classification of all possible destinations. When different providers define prices in terms of different classifications, comparisons are difficult.

Marketers and regulators have long recognized the importance of comparability for market competition. Nutritional information on food product labels is required to conform to rigid formats which include standardized units of measurement. As to the regulation of retail financial services, the following quotes from recent consumer protection reports are representative of the views of regulators:

> "The possibility to switch providers is essential for consumers to obtain the best deal. However, the Consumer Market Scoreboard 2009 showed that only 9% of consumers had switched current bank account during the previous two years. The causes again relate among others to difficulties to compare offers on banking services...In order to achieve the aims of comparable and comprehensible product information, the Commission approach has been, for some products and services...to promote the standardization of pre-contractual information obligations within carefully designed and tested formats..." (EC (2009), pp. 4,10)

> "When deciding whether to switch to another bank, consumers need clear, readily available information that they can understand, as well as the financial capability and desire to evaluate it. Ease of comparison will be affected by the structure of current account pricing. The ease with which consumers are able to compare current accounts is likely to affect their desire to do so and thus feed through to the competitive pressures that banks face." (OFT (2008), p. 89)

This paper develops a model of market competition under limited comparability. In our model, firms choose both how to price their product and how to *frame* pricing, so that consumers' "ease of comparison" is a function of the firms' framing decisions. Our aim is to address the following questions. What are the implications of limited comparability for the competitiveness of the market outcome? Do regulatory interventions aimed at enhancing comparability necessarily increase competitiveness? What determines the relation between the firms' pricing and framing decisions? What is the relation between comparability and the intensity of consumer switching?

Our model takes Bertrand competition as a starting point, and adds limited comparability as a new dimension. Two profit-maximizing firms facing a single consumer produce perfect substitutes at zero cost. They play a symmetric simultaneous-move game in which each firm chooses a price and a pricing structure - referred to as a *format*. The price is the actual payment the consumer makes to the chosen firm, whereas the format is the way in which the price is presented to the consumer. For example, a format can be a measurement unit in which prices are denominated. The consumer wishes to consume one unit and has a reservation value that is identical for both firms, regardless of their format decisions.

Given the firms' price and format decisions, the consumer chooses as follows. He is initially assigned to one firm at random. We interpret the consumer's initial firm assignment as a *default* option arising from previous consumption decisions. With a probability that is a function of the firms' chosen formats, the consumer is able to make a price comparison and chooses the cheaper firm. If the consumer is unable to make a price comparison, he buys from his default firm.

The set of available formats and the matrix of format-dependent comparison probabilities constitute what we call a "*comparability structure*". It describes the consumers' ability to compare formats. For most of the paper, we assume that the probability of comparing two formats is independent of which of them is employed by the default firm. We dub this property "*order independence*" because the comparison probability is independent of the order in which the consumer considers market alternatives. When all format pairs are compared with probability one, the model collapses to Bertrand competition.

The consumer's decision procedure in our model exhibits prudence, or "inertia", because the default is chosen whenever a comparison cannot be made. This feature of the model has some experimental support (e.g., Iyengar and Lepper (2000) and Choi, Laibson and Madrian (2009)) and chimes with the above-cited consumer protection reports, which emphasize inertia driven by limited comparability as a major cause of low switching rates and weak competition in some industries.

Equilibrium analysis hinges on whether it is possible for a firm, by an appropriate choice of formats, to nullify the relevance of format choice for the rival. When a comparability structure has this feature, we say that it satisfies "*weighted regularity*". Specifically, an order-independent comparability structure is weighted-regular if there *exists* a probability distribution over the set of available formats such that, if one firm randomizes over formats according to this distribution, the probability of a price comparison is independent of the other firm's format choice. For example, suppose that prices can only be denominated in two currencies: dollars or pounds. Assume further that the consumer makes a price comparison if and only if both firms denominate their prices in the same currency. Then, if one firm chooses to price its product in dollars or pounds with probability $\frac{1}{2}$ each, the comparison probability is $\frac{1}{2}$, regardless of the other firm's currency choice. This comparability structure thus satisfies weighted regularity.

Our yardstick for the competitiveness of the market outcome is equilibrium profits. In particular, we ask whether market forces drive profits down to the minimal level that firms can secure given the comparability structure, namely the *max-min* profit.[1] Our main result is that, in any symmetric equilibrium, *firms earn max-min profits if and only if the comparability structure is weighted-regular.* In the limit case of Bertrand competition, weighted regularity trivially holds (the comparison probability is always one), max-min profits are equal to zero and obtained in equilibrium. When the consumers' ability to make comparisons is limited, firms can typically secure positive profits, so that the market outcome cannot be strictly competitive. However, a market outcome that induces max-min profits is competitive in a "constrained", second-best sense: no extra rents are gained in addition to the minimum caused by the consumer's bounded rationality. Our result establishes a tight link between the emergence of a "constrained competitive" outcome and the comparability structure's

[1]The max-min payoff is the maximal profit that one firm can attain under the hypothesis that for any strategy it chooses, its payoff is minimized by the opponent's choice of strategy.

potential "format neutrality".

This result is also useful for deriving an equilibrium characterization for specific classes of comparability structures. For instance, "*bi-symmetric*" structures partition the set of formats into two categories, so that comparability depends only on the category to which formats belong. We obtain a closed-form characterization of the (unique) symmetric Nash equilibrium. We use this characterization to convey two important lessons: first, regulatory interventions that enhance comparability may lead to a *less* competitive market outcome; second, the relationship between comparability and consumer switching is subtle.

Throughout the paper, we follow a complexity-based interpretation of the comparability structure, and interpret formats as ways of describing prices, so that a comparability structure measures the "ease of comparison" between price formats. However, we could interpret formats more broadly, as any utility-irrelevant aspect of the product's description that affects the consumers' propensity to form a preference: an advertising message, a package design, a positioning strategy, and so forth. Saying that two formats are comparable could mean that one creates mental associations that eventually lead the consumer to pay attention to the other. This suggests a novel approach to the phenomenon of product differentiation. Conventional models regard product differentiation as a market response to consumers' heterogeneous tastes. In contrast, our approach suggests that differentiation is partly in utility-irrelevant formats, the sole function of which is to lower the probability of a price comparison. We compare these two approaches in the concluding section.

Related literature

Our paper joins recent attempts to formalize "frame-sensitive" choice. Salant and Rubinstein (2008) introduce the notion of an extended choice problem, consisting of a choice set and a frame. A choice function assigns an element in the choice set to every extended choice problem. Salant and Rubinstein conduct a choice-theoretic analysis of such extended choice functions, and identify conditions under which extended choice functions are consistent with utility maximization.[2] Ahn and Ergin (2010) axiomatize frame-dependent preferences over acts, where a frame is defined as a partition over the state space and the act is required to be measurable with respect to that partition. In contrast to these works, we focus on market implications of frame dependence rather than on axiomatic decision-theoretic analysis. Also, in our model framing creates preference incompleteness but never leads to preference reversal.

[2] Bernheim and Rangel (2009) use a similar framework to extend standard welfare analysis to frame-sensitive choices.

Eliaz and Spiegler (2011) formalize the notion that marketing activities influence the set of alternatives that consumers subject to a preference ranking. Two major features distinguish this work from our paper. First, Eliaz and Spiegler mostly interpret a frame in terms of advertising content, and assume that the consumer's propensity to consider a new market alternative is a function of its frame and the default option's *payoff-relevant* details. Second, Eliaz and Spiegler ignore price setting and assume that framing decisions are costly. The resulting market game is substantially different from ours, emphasizing the firms' trade-off between increasing their market share and lowering their advertising costs.[3]

Our paper contributes to a growing theoretical literature on the market interaction between profit-maximizing firms and boundedly rational consumers. Spiegler (2011) provides a textbook treatment of the subject. Within this literature, Spiegler (2006) and Gabaix and Laibson (2006) share the present paper's preoccupation with firms' strategic use of "confusing" pricing schemes to enhance consumers' decision errors. In Spiegler (2006), obfuscation is modeled as the introduction of noise, whereas in Gabaix and Laibson (2006) it is modeled as the shrouding of product attributes. Other papers (Ellison and Wolitzky (2008) and Wilson (2010)) stay closer to the rational-consumer paradigm, and model obfuscation as a deliberate attempt to increase consumers' search costs.

Two papers in this literature study many-firms models of price competition with boundedly rational consumers that collapse to special cases of our model in the two-firm case. Chioveanu and Zhou (2010) analyze a variant on our model in which the comparability structure is a special case of our "bi-symmetric" class, consumers are not initially assigned to default firms, and a firm is eliminated from the consumer's consideration set whenever it is comparable to a cheaper firm. They show that the market equilibrium need not converge to the competitive outcome as the number of firms tends to infinity. Carlin (2009) analyzes a model in which firms choose prices as well as a scalar variable interpreted as "complexity", and the fraction of consumers who are able to identify the cheapest firm is a separable, decreasing function of the firms' complexity decisions. He shows that, as the number of competitor increases, the equilibrium amount of complexity increases, which may cause prices to increase as well.

Our model can be viewed as an extension of a well-known model of price competition due to Varian (1980), in which consumers are divided into two groups: those

[3] Masaltioglu and Nakajima (2010) axiomatize a larger family of choice functions that involve such a formation of consisderation sets.

who make perfect price comparisons, and those who are "loyal" to the firm they are initially assigned to and thus make no comparison with other market alternatives. In Varian's model, the fraction of "loyal" consumers is exogenous, whereas in our model it is a function of the firms' endogenous format decisions. Indeed, if the comparison probability of formats is constant, format decisions are entirely irrelevant: the probability that the consumer will make a price comparison is independent of the firms' format choices, and so our model collapses into Varian's.

2 The Model

A market consists of two identical, expected-profit maximizing firms and one consumer. The firms produce a homogenous product at zero cost. The consumer wishes to buy one unit of the product. His willingness to pay for the product is 1. The two firms play a simultaneous-move game with complete information. A pure strategy for firm i is a pair (x_i, p_i), where $x_i \in X$ is a *format* and $p_i \in [0, 1]$ is the *price* of the product. A format is a way of describing the price. The set X is finite.

Given a profile $(x_i, p_i)_{i=1,2}$ of the firms' strategies, the consumer chooses a firm according to a procedure based on one primitive: a function $\pi : X \times X \to [0, 1]$ that measures the comparability of formats. Specifically, $\pi(x, y)$ is the probability that the consumer is able to compare the format y to the format x. The pair (X, π) is called a *comparability structure* (CS henceforth). Throughout the paper, we assume that $\pi(x, x) = 1$, that is, each format is perfectly comparable to itself.[4]

The choice procedure is as follows. The consumer is randomly assigned to a firm, say firm i, each firm being selected with equal probability. We interpret this assignment as the consumer's *default*. With probability $\pi(x_i, x_j)$, the consumer makes a price comparison, and switches from firm i to firm $j \neq i$ if $p_j < p_i$. Otherwise, he chooses the default option. Thus, the consumer stays with the default firm in two cases: either no price comparison is made, or it is made and the consumer realizes that the default is not more expensive than the alternative.

When $\pi(x, y) = \pi(y, x)$ for all $x, y \in X$, we say that (X, π) is *order-independent*. When the CS is order-independent, the probability that the consumer buys from any given firm is independent of the identity of the firm to which he is initially assigned (hence the term "order independence"). When $\pi(x, y)$ is either 0 or 1 for every $x, y \in X$, we say that (X, π) is *deterministic*. Non-deterministic structures allow for heterogeneity

[4]This assumption is made for expositional simplicity. All our results continue to hold - subject to minor adjustments in Section 4 - if we assume instead that $\pi(x, x) > 0$ for all $x \in X$.

in the population from which our consumer is drawn, in that $\pi(x, y)$ can be interpreted as the fraction of consumers in this population who are able to compare y to x.

This procedure generates the following payoff function for the firms:

	$p_1 < p_2$	$p_1 > p_2$	$p_1 = p_2$
firm 1	$p_1 \frac{1+\pi(x_2,x_1)}{2}$	$p_1 \frac{1-\pi(x_1,x_2)}{2}$	$\frac{p_1}{2}$
firm 2	$p_2 \frac{1-\pi(x_2,x_1)}{2}$	$p_2 \frac{1+\pi(x_1,x_2)}{2}$	$\frac{p_2}{2}$

For example, let $X = \{a, b\}$, $\pi(a, b) = \pi(b, a) = q$. Suppose that $p_1 < p_2$. If $x_1 = x_2$, firm 1 earns p_1 and firm 2 earns zero, because by assumption, $\pi(x_1, x_2) = 1$. If $x_1 \neq x_2$, firm 1 earns $p_1 \cdot \frac{1+q}{2}$ and firm 2 earns $p_2 \cdot \frac{1-q}{2}$.

Remarks on the consumer's choice procedure

Formats represent the language that firms choose to describe their prices. They are utility-irrelevant: the consumer's willingness to pay for a product is independent of its format. For instance, a format could be a particular unit of volume (e.g., metric or British) per which prices are displayed. In this case, π represents the consumer's ability to convert one unit into another. The assumption of order independence appears to be a good approximation here: converting one measurement unit to another should be equally hard to perform in either direction. When the product is a savings plan, a format can be a time unit for which the implicit interest rate is defined. When the product is a mortgage, a format can be a particular way of separating a given cash flow into principal and interest payments. In these cases, π represents the consumer's financial numeracy (e.g., his understanding of compounding - see Banks and Oldfield (2007) for a related empirical study).

The assumption that comparability of market alternatives depends only on their formats, and not on the actual prices, is quite strong. Since the modeler can always incorporate prices into the definition of formats, the real assumption made here is that a firm's choice of format does not restrict the set of prices it can charge. This assumption clearly entails a loss of generality. Suppose, for example, that firms sell a product with attributes A and B; a format is a price pair (p_A, p_B), and the price paid by the consumer is $p_A + p_B$. Then, a firm's choice of format uniquely determines its price, contrary to our assumption.

Firms in our model cannot use their format decisions to fool the consumer into paying a price above the reservation value 1, even when he is unable to compare formats. One could argue that if consumers have limited ability to understand the prices they are facing, firms should be able to charge prices above their willingness to pay. One justification for our assumption is the possible presence of an implicit ex-post participation constraint: once the consumer realizes he has been fooled, he can opt out without paying anything. The presence of a default option makes this justification particularly sensible. Even if a consumer does not understand the price structure of his default option, he can appreciate whether he actually pays more than his reservation value; and he will never pay in excess of his reservation value for the alternative, because the choice procedure implies that he can switch to it only after comparing its price to the default.

Graph representations

It will often be convenient to represent CSs as random directed *graphs*, where X is the set of nodes and $\pi(x,y)$ is the probability of a directed link from x to y. Order-independent CSs will be represented as random *non-directed* graphs, where $\pi(x,y) = \pi(y,x)$ is the probability of a link between x and y. A graph representation entails no loss of generality: its main role is to visualize CSs that involve many formats and suggest fruitful notions of comparability.

Mixed strategies

Mixed strategies will play an important role in the analysis. A mixed strategy is a joint probability measure over all feasible pairs (x,p). Let $\Delta(X)$ denote the set of probability distributions over X. We will usually represent mixed strategies as a pair $(\lambda,(F^x)_{x\in X})$, where $\lambda \in \Delta(X)$ is referred to as the (*marginal*) *format strategy*, and F^x is the pricing *cdf* conditional on the format x.[5] The marginal pricing *cdf* induced by this mixed strategy is then $F \equiv \sum_{x\in X} \lambda(x)F^x$. Given a probability distribution ϕ, *Supp* (ϕ) denotes its support. For every interval $I = [a,b]$, $a < b$, in *Supp*(F), let λ^I denote the format strategy conditional on the event that the price realization lies in I. Given a *cdf* F^y, let F^{y-} denote its left limit - that is, $F^{y-}(p) \equiv \lim_{p'\to p^-} F^y(p')$.

When firm j plays the mixed strategy $\left(\lambda_j,(F_j^x)_{x\in X}\right)$, firm i's expected payoff from

[5] Naturally, the specification of a conditional pricing *cdf* for formats outside the support of λ is redundant and made solely for ease of notation.

the pure strategy (x_i, p_i) is

$$\frac{p_i}{2}\left(1+\sum_{y\in X}\lambda_j(y)\left[(1-F_j^y(p_i))\cdot\pi(y,x_i)-F_j^{y-}(p_i)\cdot\pi(x_i,y)\right]\right)$$

This expression makes clear that a firm's market share consists of its initial clientele (namely, 50% of the market), plus consumers it is able to attract when the rival charges a higher price, minus consumers it loses when the rival charges a lower price.

2.1 Equilibrium

We will focus on symmetric Nash equilibria in the game played by the two firms. It can be shown that a *symmetric* mixed-strategy equilibrium must exist. The proof is a direct application of Corollary 5.3 in Reny (1999), and is omitted. Throughout the paper, $(\lambda, (F^x)_{x\in X})$ denotes an equilibrium strategy in a symmetric Nash equilibrium. By standard arguments, the assumption that $\pi(x,x) > 0$ for all $x \in X$ implies that F^x is continuous for any $x \in Supp(\lambda)$, and that the marginal pricing strategy F is continuous and strictly increasing. In addition, there exists $p^l \in (0,1)$ such that $Supp(F) = [p^l, 1]$. These properties are entirely conventional and their proofs are therefore omitted.

2.2 Hide and Seek

Our analysis will make use of an auxiliary two-player, zero-sum game, associated with each CS (X, π), which is a generalization of familiar games such as Matching Pennies. The players (not to be identified with the firms), named *hider* and *seeker* and denoted by h and s, share the same action space X. Given the action profile (x_h, x_s), the hider's payoff is $-\pi(x_h, x_s)$ and the seeker's payoff is $\pi(x_h, x_s)$. This game will be referred to as the *hide-and-seek* game associated with (X, π).

Given a mixed-strategy profile (λ_h, λ_s) in this game, the seeker's payoff is the probability that he finds the hider - that is,

$$v(\lambda_h, \lambda_s) = \sum_{x\in X}\sum_{y\in X}\lambda_h(x)\,\lambda_s(y)\,\pi(x,y)$$

By the fundamental Minimax Theorem, since the hide-and-seek game is a finite zero-sum game, it has a value

$$v^* = \max_{\lambda_s} \min_{\lambda_h} v(\lambda_h, \lambda_s) = \min_{\lambda_h} \max_{\lambda_s} v(\lambda_h, \lambda_s)$$

which is equal to the seeker's payoff in all Nash equilibria of the game.

To see the relevance of this auxiliary game to our model, suppose that firm 1's marginal format and pricing strategies are λ and F, respectively, where the latter is continuous with support $[p^l, p^u]$. When firm 2 charges the price p^l (p^u), it selects a format that maximizes (minimizes) the probability of a price comparison. Hence, it behaves as a seeker (hider) in the hide-and-seek game, facing a hider (seeker) who plays λ. Since firm 2 can play a mixed format strategy that enforces a comparison probability of at least (most) v^*, it can secure a market share of at least $\frac{1}{2} + \frac{1}{2}v^*$ $(\frac{1}{2} - \frac{1}{2}v^*)$. This is a lower bound on the market share that a firm obtains in any Nash equilibrium, when its price realization is the lowest (highest) in the equilibrium distribution.

The hide-and-seek game is also useful for calculating the max-min payoff in our model. Max-min payoffs represent the lowest possible profit consistent with firms' individual rationality and consumers' bounded rationality. Therefore, when firms earn the max-min payoff, the market outcome can be viewed as "constrained competitive". To calculate the max-min payoff, note that the worst-case scenario for a firm is that its opponent plays $p = 0$ and adopts the seeker's max-min format strategy. A best-reply is then to play $p = 1$ with a format strategy that minimizes the probability of a price comparison. Since the minimum probability is v^*, the max-min payoff is $\frac{1}{2}(1 - v^*)$.

2.3 When is the Equilibrium Outcome Competitive?

The above discussion makes it obvious that, when $v^* < 1$, the max-min payoff is strictly positive and firms necessarily earn strictly positive profits in any Nash equilibrium. It is sensible to ask whether comparability structures that generate competitive equilibrium outcomes can be characterized. A simple necessary and sufficient condition, namely the existence of a "universally comparable" format, provides a natural and intuitive characterization for symmetric equilibria.

Proposition 1 *In any symmetric Nash equilibrium, both firms play $p = 0$ with probability one if and only if there exists a format $x^* \in X$ such that $\pi(y, x^*) = 1$ for every $y \in X$.*

The condition for a competitive symmetric equilibrium outcome is also necessary and sufficient for the max-min payoff to be zero. Thus, symmetric equilibrium profits are equal to zero if and only if the max-min payoff is zero.

For an intuition behind the sufficiency of the condition for a competitive equilibrium outcome, suppose that firms charge positive prices in equilibrium. Consider the *lowest* price p^* in the support of the marginal pricing strategy, at which firms do *not* adopt a universally comparable format. Suppose that one firm switches to such a format at p^*. By the definition of p^*, the comparison probability is unaffected if the rival charges a lower price, but increases if the rival charges a higher price. The deviation raises the firm's market share, hence it is profitable. It follows that in equilibrium firms use only universally comparable formats, and standard Bertrand competition arguments apply.

For the rest of the paper, we will rule out competitive outcomes by assuming the absence of a universally comparable format:

Condition 1 *For every $x \in X$ there exists $y \in X$ such that $\pi(y, x) < 1$.*

Note that, under this condition, any Nash equilibrium must be *mixed*. To see why, assume that firm i plays a pure strategy (x_i, p_i). If $p_1 = p_2 = p > 0$, then any firm i can profitably deviate to the strategy $(x_j, p - \varepsilon)$, where $\varepsilon > 0$ is arbitrarily small, and raise its payoff from $\frac{1}{2}p$ to $p - \varepsilon$. If $p_i = 0$, firm i earns zero profits, contradicting the observation that the firms' max-min payoffs are strictly positive. Finally, if $p_i < p_j$, firm i can profitably deviate to (x_j, p'), where p' is a price between p_i and p_j.

3 Equilibrium Analysis under Order Independence

In this section, we restrict attention to *order-independent* CSs, and address the question whether symmetric equilibrium profits are equal to or above the max-min level - that is, whether market forces induce a "constrained competitive" outcome. As we shall see, the answer to this question carries rich implications for the structure of equilibria.

To illustrate a symmetric equilibrium, consider the following deterministic CS. Let $|X| = m > 1$, and assume that $\pi(x, y) = 0$ whenever $x \neq y$. Suppose that each firm mixes over X according to the uniform distribution, *independently* of the price it charges. Given that firm i plays this format strategy, the probability that the consumer makes a comparison is $\frac{1}{m}$, whatever firm $j \neq i$ does. Since the conditional pricing *cdf*'s are the same as the marginal pricing *cdf* F, charging a price p in the support of F, it earns an expected profit of

$$p \cdot \frac{1}{2}[1 + \frac{1}{m}(1 - F(p)) - \frac{1}{m}F(p)]$$

By construction, $F(p^l) = 0$ and $F(1) = 1$. In a mixed-strategy equilibrium, every $p \in [p^l, 1]$ must generate the same profit. It follows that the equilibrium payoff is $\frac{m-1}{2m}$, which is equal to the max-min. The marginal pricing *cdf* is

$$F(p) = 1 - \frac{m-1}{2} \cdot \frac{1-p}{p} \tag{1}$$

with

$$p^l = \frac{m-1}{m+1}$$

Two features of this equilibrium are noteworthy. First, firms earn max-min payoffs. Second, price and format decisions are statistically independent. In particular, the comparison probability is the same for any of the firms' price realizations.

3.1 Weighted Regularity

A notable feature of the above example is the existence of a format strategy for one firm which induces the same comparison probability for any format choice of the other firm. This notion of "uniform comparability" across formats plays a central role in our equilibrium analysis.

Definition 1 *An order-independent comparability structure* (X, π) *is weighted-regular if there exist* $\beta \in \Delta(X)$ *and* $\bar{v} \in [0,1]$ *such that, for any* $x \in X$,

$$\sum_{y \in X} \beta(y)\, \pi(x, y) = \bar{v}$$

We say that such β *verifies weighted regularity.*

The economic interpretation of weighted regularity is that it is possible for one firm to make its opponent indifferent among all formats - in other words, to "neutralize" the relevance of framing for the rival firm's competitive strategy.

Weighted regularity generalizes the familiar notion of *regular graphs.* A CS (X, π) is regular if there exists a number $\bar{v} > 0$ such that $\sum_{y \in X} \pi(x, y) = \bar{v}$ for any $x \in X$. In this case, the uniform distribution over X verifies weighted regularity. Here are some additional examples:

1. Consider a deterministic CS that partitions X into m classes, such that $\pi(x,y) = 1$ if and only if x and y belong to the same class. (The CS in the example that opened this section falls into this category.) Any distribution over formats that assigns probability $\frac{1}{m}$ to each equivalence class verifies weighted regularity, with $\bar{v} = \frac{1}{m}$.

2. Let $X = \{1, 2, ..., n\}$, where n is even. Assume that for every distinct $x, y \in X$, $\pi(x,y) = \frac{1}{2}$ if $|y - x| \in \{1, n-1\}$, and $\pi(x,y) = 0$ otherwise. Among others, a uniform distribution over all odd-numbered formats verifies weighted regularity, with $\bar{v} = \frac{2}{n}$.

The following lemma establishes an equivalent definition of weighted regularity, which makes use of the auxiliary hide-and-seek game of Section 2.2.

Lemma 1 *In an order-independent CS (X, π), the distribution $\lambda \in \Delta(X)$ verifies weighted regularity if and only if (λ, λ) is a Nash equilibrium in the associated hide-and-seek game.*

Proof. (Only if) Suppose that λ verifies weighted regularity. If one of the players in the associated hide-and-seek game plays λ, every strategy for the opponent - including λ itself - is a best-reply. Therefore, λ is a symmetric equilibrium strategy.
(If) Suppose that (λ, λ) is a Nash equilibrium in the associated hide-and-seek game. Denote $v(\lambda, \lambda) = \bar{v}$. If some format attains a higher (lower) probability of a price comparison than $\bar{v}$, then λ cannot be a best-reply for the seeker (hider). Therefore, every format generates the same probability of a price comparison - namely $\bar{v}$ - against λ. ■

Thus, weighted regularity implies that the format strategy of a firm that maximizes comparability need not be distinct from the format strategy of a firm that minimizes it.

An immediate corollary of Lemma 1 is that when weighted regularity is satisfied, the induced comparison probability $\bar{v}$ is in fact the value of the hide-and-seek game, v^*. In particular, all format strategies that verify weighted regularity for a given CS generate the same comparison probability.

3.2 The Main Results

We are now ready for the two main results of this paper. First, we establish equivalence between weighted regularity and the property that firms earn max-min payoffs in a symmetric equilibrium.

Theorem 1 *In any symmetric equilibrium, firms earn max-min payoffs if and only if (X, π) is weighted-regular. Furthermore, if (X, π) is weighted-regular, then in any symmetric equilibrium each firm's marginal format strategy verifies weighted regularity.*

Proof. We first show that, under weighted regularity, in any symmetric equilibrium firms earn max-min payoffs and each firm's marginal format strategy verifies weighted regularity. Fix a symmetric Nash equilibrium. For every $p \in [0,1]$, define $s(p)$ as a firm's maximum market share conditional on charging the price p, namely

$$s(p) = \max_{x \in X} \left\{ \frac{1}{2} \cdot \left(1 + \sum_{y \in X} \lambda(y) \left[(1 - F^y(p)) \cdot \pi(y,x) - F^y(p) \cdot \pi(x,y) \right] \right) \right\}.$$

We will show that

$$s(p) = \frac{1}{2}[1 + (1 - F(p))v^* - F(p)v^*] \tag{2}$$

for every $p \in [p^l, 1]$.

First notice that, since $F(p)$ is continuous, $s(p)$ is also continuous. By weighted regularity, each firm can enforce a constant comparison probability v^*, independently of the opponent's action, and thus obtain a market share equal to the R.H.S of (2). Thus

$$\int_{p^l}^{1} s(p)dF(p) \geq \frac{1}{2} \int_{p^l}^{1} [1 + (1 - F(p))v^* - F(p)v^*]dF(p) \tag{3}$$

The R.H.S of this inequality is equal to

$$\frac{1}{2} + \frac{1}{2}v^* - v^* \int_{p^l}^{1} F(p)dF(p)dp = \frac{1}{2} \tag{4}$$

By profit maximization, each firm's ex-ante equilibrium market share must be equal to the L.H.S of (3). By the equilibrium symmetry, each firm's ex-ante market share is equal to $\frac{1}{2}$, and thus (2) follows. Since $s(1) = \frac{1}{2}(1 - v^*)$, firms earn max-min payoffs. Also, $s(p^l) = \frac{1}{2}(1 + v^*)$. Since $s(p^l) = \frac{1}{2}[1 + \max v(\cdot, \lambda)]$ and $s(1) = \frac{1}{2}[1 - \min v(\cdot, \lambda)]$, it follows that $\max v(\cdot, \lambda) = \min v(\cdot, \lambda) = v^*$, and hence λ verifies weighted regularity.

We now show that max-min payoffs imply weighted regularity. Assume that firms earn max-min payoffs in some symmetric equilibrium. Then, $s(1) = \frac{1}{2}(1 - v^*)$. Recall that $s(p^l) \geq \frac{1}{2}(1+v^*)$. Assume that the inequality is strict. Since each firm's ex-ante market share is $\frac{1}{2}$, it follows from (4) that there exists a price $p \in (p^l, 1)$ such that

$$s(p) < \frac{1}{2}[1 + (1 - F(p))v^* - F(p)v^*]$$

Since no format strategy can yield a market share above the equilibrium one at a given price, the following inequality holds for any $\theta \in \Delta(X)$

$$\sum_{x\in X}\sum_{y\in X}\theta(x)\lambda(y)\pi(x,y)(1 - 2F^y(p)) < (1 - F(p))v^* - F(p)v^* \tag{5}$$

which can be rewritten as

$$v(\theta, \lambda) - 2\sum_{x\in X}\sum_{y\in X}\theta(x)\lambda(y)\pi(x,y)F^y(p) < v^*(1 - 2F(p))$$

Because firms earn max-min payoffs by hypothesis, λ max-minimizes v, and hence $v(\theta, \lambda) \geq \max\min(v) = v^*$. Then, it follows that

$$\frac{\sum_{x\in X}\sum_{y\in X}\theta(x)\lambda(y)\pi(x,y)F^y(p)}{F(p)} = v(\theta, \lambda^{[p^l,p]}) > v^*$$

for every θ, contradicting the fact that $v^* = \max\min(v)$. Therefore, $s(p^l) = \frac{1}{2}(1+v^*)$. This implies that $\min v(\cdot, \lambda) = \max v(\cdot, \lambda) = v^*$, hence (X, π) is weighted-regular. ■

The economic significance of this result is that it establishes a tight link between two aspects of market interaction. On one hand, max-min equilibrium profits occur when market forces have driven industry profits to the "constrained competitive" level - i.e., the lowest profit compatible with consumers' bounded rationality and firms' individual rationality. On the other hand, weighted regularity is present when the effect of framing on price comparisons is potentially neutralized. The theorem states that the two properties are equivalent: a constrained competitive market outcome emerges in equilibrium if and only if framing is potentially neutralized.

For an approximate intuition for Theorem 1, recall that when firms earn max-min payoffs in a symmetric equilibrium, their marginal format strategy max-minimizes the probability of a price comparison - namely, it is a max-min strategy for the seeker in the associated hide-and-seek game. Also recall that when a firm charges a price toward the high (low) end of the price distribution, it has an incentive to select a format as if it were the hider (seeker) in the hide-and-seek game. When the CS is not weighted-regular, “acting like a hider” is necessarily distinct from “acting like a seeker”. Therefore, since the marginal format strategy averages out the firms' choices of formats across all prices, it is impossible for the marginal format strategy to coincide with a seeker's max-min strategy in the hide-and-seek game. As a result, the firms' equilibrium payoff exceeds the max-min level.

In contrast, when the CS is weighted-regular, a firm can choose to play a format strategy that acts “like a hider” and “like a seeker” at the same time. If a firm did not know its relative position in the price distribution, it could use this strategy as a “hedge” to secure a comparison probability of exactly v^*. Thus, the R.H.S of (2) is a lower bound on the firm's market share for any price p it charges in equilibrium. Since each firm gets a market share of 50% ex-ante, this bound is binding in equilibrium, which implies that, when a firm charges $p = 1$, it earns the max-min payoff $\frac{1}{2}(1 - v^*)$.

Theorem 1 has an immediate implication for the structure of the firms' marginal pricing strategy under weighted regularity.

Corollary 1 *Suppose that (X, π) is weighted-regular. Then, in any symmetric equilibrium, firms play a marginal pricing strategy given by the cdf*

$$F^*(p) = 1 - \frac{1 - v^*}{2v^*} \cdot \frac{1 - p}{p} \tag{6}$$

defined over the support $[\frac{1 - v^*}{1 + v^*}, 1]$.

Proof. The proof of Theorem 1 establishes that the firms' symmetric equilibrium market share as a function of the price they charge is given by (2). We also established that in equilibrium each firm earns the max-min payoff $\frac{1}{2}(1 - v^*)$. Therefore, every price p in the support of the equilibrium marginal pricing *cdf* F generates a payoff of

$$\frac{p}{2}[1 + (1 - F(p))v^* - F(p)v^*] = \frac{1}{2}(1 - v^*)$$

Given the support $[p^l, 1]$, by definition $F(p^l) = 0$ and $F(1) = 1$. The unique solution to this functional equation is F^*. ■

It should be noticed that the pricing strategy defined in Equation (6) of Corollary 1 is also the equilibrium strategy in the two-firm case in Varian (1980). The intuition for this is simple. Under weighted regularity the firms' equilibrium market share is determined as a function of the prices they charge given a constant comparison probability (set at v^*), as in Varian (1980).

Our second main result establishes the relation between the probability with which consumers make a price comparison and the realizations of the firms' prices. We will say that a symmetric equilibrium exhibits a *constant comparison probability* if $v(\lambda^I, \lambda^J)$ is the same for every pair of closed intervals $I, J \subseteq [p^l, 1]$.

Theorem 2 *A symmetric equilibrium exhibits a constant comparison probability if and only if (X, π) is weighted-regular. Furthermore, if (X, π) is weighted-regular, the constant equilibrium comparison probability is v^*.*

Proof. First, suppose that the comparison probability is constant and λ does not verify weighted regularity, that is, there exist $x, y \in X$ such that $v(x, \lambda) < v(y, \lambda)$. Since F is strictly increasing and continuous, it can be easily shown that there exist $\hat{p}$ close to p^l and $\tilde{p}$ close to 1 such that $v(\lambda^{[p^l,\hat{p}]}, \lambda) > v(\lambda^{[\tilde{p},1]}, \lambda)$, that is, the comparison probability is not constant, a contradiction.

Now suppose that (X, π) is weighted-regular. Recall equation (2) in the proof of Theorem 1. This equation can be written as

$$\max_\theta [v(\theta, \lambda) - 2v(\theta, \lambda^{[p^l,p]})F(p)] = \max_\theta [2v(\theta, \lambda^{[p,1]})(1 - F(p)) - v(\theta, \lambda)] = v^*(1 - 2F(p))$$

for every $p \in [p^l, 1]$. Since λ verifies weighted regularity, $v(\theta, \lambda) = v^*$ for every θ. Therefore:

$$\min_\theta v(\theta, \lambda^{[p^l,p]}) = \max_\theta v(\theta, \lambda^{[p,1]}) = v^*$$

Hence, for every $p \in [p^l, 1]$, $\lambda^{[p^l,p]} \in \arg\max \min v$ and $\lambda^{[p,1]} \in \arg\min \max v$. Thus, for every $p, q \in [p^l, 1]$, $(\lambda^{[q,1]}, \lambda^{[p^l,p]})$ is a Nash equilibrium in the hide-and-seek game, and therefore $v(\lambda^{[q,1]}, \lambda^{[p^l,p]}) = v^*$. Now consider two arbitrary price intervals $[a, b], [c, d] \subseteq [p^l, 1]$. We established that $v(\lambda^I, \lambda^J) = v^*$ for every $I \in \{[p^l, a], [p^l, b]\}$ and every $J \in \{[c, 1], [d, 1]\}$. It follows that $v(\lambda^{[a,b]}, \lambda^{[c,d]}) = v^*$. ■

The proof of Theorem 1 establishes that under weighted regularity, a firm's market share when it charges a price p is $\frac{1}{2}[1 + (1 - F(p))v^* - F(p)v^*]$ - that is, it is *as if* the firm faces a constant comparison probability of v^*. Theorem 2 shows that this is not merely an "as if" property; it actually holds in symmetric equilibrium if and only if the CS is weighted-regular.

The previous results shed some light on whether a firm's pricing and format decisions exhibit correlation. An immediate corollary of Theorem 2 is that, when weighted regularity is violated, price and format decisions *must* be correlated. The reason is simple: if these decisions are statistically independent, it follows that each firm adopts the same format strategy when it charges the highest or the lowest price. But since such format strategies respectively minimize and maximize $v(\cdot, \lambda)$, λ verifies weighted regularity, a contradiction.

The converse, however, is not true: weighted regularity does *not* rule out correlation between firms' equilibrium price and format decisions. A trivial example is obtained taking a weighted-regular CS and replicating one of its formats, so that the new CS contains two distinct formats x, x' with $\pi(x, y) = \pi(x', y)$ for every $y \in X$. In this case, one can easily construct an equilibrium in which the format x is associated with low prices while the format x' is associated with high prices. For a non-trivial example, consider the deterministic nine-node graph given by Figure 1. A uniform distribution over the six bold-face nodes verifies weighted regularity ($\bar{v} = \frac{1}{3}$). By Theorem 1, this is the marginal format strategy in any symmetric equilibrium. However, one can construct an equilibrium in which price and format decisions are correlated. Specifically, the three peripheral formats are played with probability $\frac{1}{3}$ each conditional on $p \in [\frac{2}{3}, 1]$, while their internal neighbors are played with probability $\frac{1}{3}$ each conditional on $p \in [\frac{1}{2}, \frac{2}{3})$.[6]

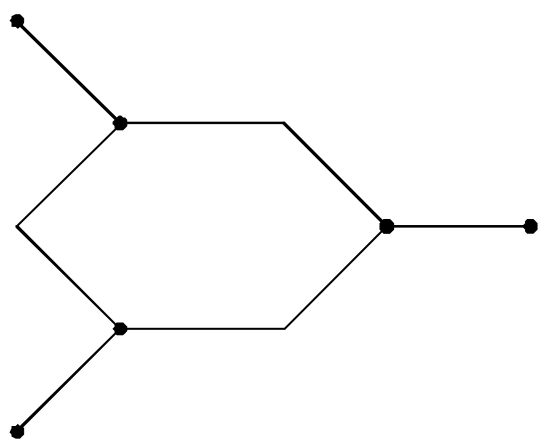

(Figure 1)

It should be noted that the associated hide-and-seek game in this example has an asymmetric equilibrium in which the three peripheral formats are played with probability $\frac{1}{3}$ each by the hider, while their internal neighbors are played with probability $\frac{1}{3}$

[6] This example also illustrates that weighted regularity does *not* imply that in equilibrium, firms are indifferent among all formats at *all* prices. For example, when a firm charges the cutoff price $p = \frac{2}{3}$, it strictly prefers the bold-face nodes to any of the three other nodes. The indifference among all formats holds at the extreme prices $p = \frac{1}{2}$ and $p = 1$.

each by the seeker. In contrast, when a CS is weighted-regular and the hide-and-seek game has a unique Nash equilibrium (which is therefore symmetric), there is a unique symmetric equilibrium in our model, and in this equilibrium the firms' price and format decisions must be independent. This result follows immediately from the proof of Theorem 2. Under weighted regularity, for every firm i and every price p in the support of the equilibrium strategy, $\lambda^{[p^l,p]}$ max-minimizes v and $\lambda^{[p,1]}$ min-maximizes v. Thus, $\lambda^{[p^l,p]} = \lambda^{[p,1]} = \lambda$, where λ is the unique equilibrium strategy in the hide-and-seek game, and thus each firm plays λ independently of the price it charges.

Recall the example of m non-comparable formats examined at the beginning of this section. Since the hide-and-seek game associated with this CS has a unique Nash equilibrium, it follows that the equilibrium we constructed for this structure is the *unique* symmetric equilibrium. This observation demonstrates the usefulness of Theorem 1 in obtaining strong equilibrium characterizations.

4 Bi-Symmetric Comparability Structures

In this section, we provide a complete characterization of symmetric equilibria for a specific family of intuitive and easily interpretable CSs. Comparative statics will show that regulatory interventions that improve comparability in a "local" sense can have negative effects on the competitiveness of the market outcome, as well as on the amount of consumer switching.

An order-independent comparability structure (X, π) is *bi-symmetric* if X can be partitioned into two sets, W and Z, such that for every *distinct* $x, y \in X$:

$$\pi(x,y) = \begin{cases} q_W & if \quad x,y \in W \\ q_Z & if \quad x,y \in Z \\ q & if \quad x \in W,\, y \in Z \end{cases}$$

where $\max\{q_W, q_Z, q\} < 1$. Thus, a bi-symmetric CS is fully determined by five parameters: the number of formats in the two categories, Z and W, the comparability of different formats within each category and the comparability of the two categories.

Let n_I denote the number of formats in category $I \in \{W, Z\}$. Let

$$q_I^* = \frac{1 + q_I \cdot (n_I - 1)}{n_I}$$

be the "average comparability" within category $I \in \{W, Z\}$. Without loss of generality, assume $q_Z^* \geq q_W^*$.

Bi-symmetric CSs are attractive because they enable us to capture limited-comparability "stories" with simple restrictions on parameter values. When $q < \min\{q_W, q_Z\}$, we may interpret formats within any category $I \in \{W, Z\}$ as similar and therefore more comparable than formats that belong to different categories. In contrast, when $q_W < q < q_Z$, we may interpret the formats in category Z as inherently simpler than those in W (possibly because they contain translations or convertors that are absent from the formats in W).

Example: The Star Graph

The star graph in Figure 2 represents a bi-symmetric CS, with $n_Z = 1$, $n_W = 4$, and $q_W = 0$. A simple scenario for this CS is that the product traded in the market can be priced in five different currencies, one major and four minor ones. The consumer is able to compare prices denominated in different currencies only if he knows the exchange rate. With probability q, he knows the exchange rate between the major currency and any minor one. He does not know the exchange rates between minor currencies.

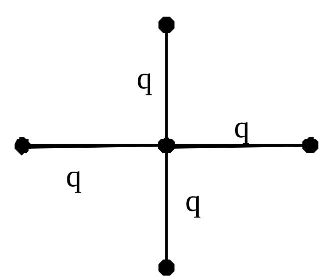

(Figure 2)

4.1 Weighted-Regular Structures

One can verify that a bi-symmetric CS is weighted-regular if and only if

$$(q_W^* - q)(q_Z^* - q) \geq 0 \tag{7}$$

When this condition holds with $q_W^* = q_Z^* = q$, there is a continuum of format strategies that verify weighted regularity, with $v^* = q$. When it holds with $q_W^* \neq q$ or $q_Z^* \neq q$, the unique format strategy λ^* that verifies weighted regularity assigns probability

$$\frac{q_W^* - q}{(q_W^* - q) + (q_Z^* - q)}$$

to the set Z, and mixes uniformly within category. The hide-and-seek game has (λ^*, λ^*) as the unique Nash equilibrium, and

$$v^* = \frac{q_W^* q_Z^* - q^2}{(q_W^* - q) + (q_Z^* - q)} \tag{8}$$

The equilibrium characterization below follows directly from our results in the previous section.

Proposition 2 *Suppose that (X, π) is bi-symmetric, and that $(q_W^* - q)(q_Z^* - q) \geq 0$, with $q_W^* \neq q$ or $q_Z^* \neq q$. Then, there is a unique symmetric Nash equilibrium, in which firms play the format strategy λ^*, and independently the pricing strategy (6), where v^* is given by (8). Firms earn the max-min payoff $\frac{1}{2}(1 - v^*)$.*[7]

Let us illustrate this result for the star graph. Condition (7) simplifies to $q \leq \frac{1}{4}$. When this holds, there is a unique symmetric equilibrium where firms play a format strategy that assigns probability $\frac{1-4q}{5-8q}$ to the core format and probability $\frac{1-q}{5-8q}$ to each of the four peripheral formats. Independently of their format choice, firms randomize over prices according to the *cdf* (6), where the constant comparison probability is $v^* = \frac{1-4q^2}{5-8q}$.

4.2 Non-Weighted-Regular Structures

Recall that $q_Z^* \geq q_W^*$. If a bi-symmetric CS violates condition (7), the value of the hide-and-seek game is $v^* = q$, since there exists a Nash equilibrium in which the seeker (hider) mixes uniformly over Z (W). We use this observation to construct a symmetric equilibrium strategy with a "cutoff" structure. The support of the marginal pricing *cdf* F is $[p^l, 1]$, $0 < p^l < 1$, and there exists a price $p^m \in (p^l, 1)$ such that the format strategy conditional on any price $p \in [p^l, p^m)$ is the *uniform distribution* over Z, and

[7] When $q_Y^* = q_Z^* = q$, the result is slightly weaker. In symmetric equilibrium, the marginal format strategy verifies weighted regularity, and the pricing strategy is (6), with $v^* = q$. However, the infinite number of format strategies that verify weighted regularity can give rise to payoff-irrelevant correlation between price and format choices.

the format strategy conditional on any price $p \in (p^m, 1]$ is the *uniform distribution* over W.

Consider again the CS in Figure 2 and assume that $q > \frac{1}{4}$, so that the condition for weighted regularity is violated. In the cutoff equilibrium, firms mix uniformly over peripheral formats conditional on charging a price above the cutoff, and play the core format conditional on charging a price below it. The intuition is that when $q > \frac{1}{4}$, the core format *dominates* peripheral formats in terms of comparability, in that it generates a higher comparison probability *regardless* of the rival firm's format decision. Therefore, when a firm charges a low (high) price, it has a clear-cut incentive to adopt a format in the core (periphery).

To construct the cutoff equilibrium in the general case, first note that the total probability that the marginal format strategy assigns to the set Z (W) is $F(p^m)$ $(1 - F(p^m))$. Conditional on charging p^m, firms should be indifferent among all formats. Therefore:

$$-F(p^m)q_Z^* + (1 - F(p^m))q = (1 - F(p^m))q_W^* - F(p^m)q$$

Rearranging, we get:

$$F(p^m) = \frac{q - q_W^*}{q_Z^* - q_W^*} \tag{9}$$

The conditional pricing strategies are given by the following pair of functional equations which represent indifference among all prices in the support of F. Let F^Z (F^W) denote the pricing *cdf* conditional on playing a format in Z (W). Setting a price $p \in [p^l, p^m]$ yields the same profits as setting $p = 1$, that is,

$$\frac{p}{2}[1 + F(p^m)(1 - 2F^Z(p))q_Z^* + (1 - F(p^m))q] = \frac{1}{2}[1 - F(p^m)q - (1 - F(p^m))q_W^*] \tag{10}$$

Similarly, for every $p \in [p^m, 1]$:

$$\frac{p}{2}[1 + (1 - F(p^m))(1 - 2F^W(p))q_W^* - F(p^m)q] = \frac{1}{2}[1 - F(p^m)q - (1 - F(p^m))q_W^*] \tag{11}$$

The R.H.S on each of these two equations represents the firms' equilibrium payoff. Plugging in (9), the equilibrium payoff is

$$\frac{1}{2}\left(\frac{q - q_W^*}{q_Z^* - q_W^*}(1 - q) + \frac{q_Z^* - q}{q_Z^* - q_W^*}(1 - q_W^*)\right) \tag{12}$$

Observe that this expression exceeds the max-min level $\frac{1}{2}(1-q)$, in accordance with Theorem 1. The following proposition establishes that there are no other symmetric equilibria.[8]

Proposition 3 *Suppose that (X, π) is bi-symmetric, and that $(q_W^* - q)(q_Z^* - q) < 0$. Then, there is a unique symmetric Nash equilibrium, which is the cutoff equilibrium characterized by (9)-(11). The firms' equilibrium payoff is given by (12).*

As mentioned earlier, the classification of bi-symmetric CSs into those that satisfy weighted regularity and those that do not corresponds to two different interpretations of the categories W and Z. The results in this section imply that when parameter values fit situations in which the categorization of formats captures their relative complexity, the firms' equilibrium strategy displays correlation between price and format decisions - simple (complex) formats are coupled with low (high) price realizations - and firms earn "collusive" profits. In contrast, when parameter values fit situations in which the categorization of formats captures their relative similarity, the equilibrium strategy displays format-price independence and firms earn max-min payoffs.

5 Does Greater Comparability Lead to a More Competitive Outcome?

The basic idea that underpins the quotes in the Introduction is that greater comparability of price formats makes a market more competitive and benefits consumers. Indeed, if consumers face a fixed set of format-price pairs, switching from a CS (X, π) to another structure (X, π') for which $\pi'(x, y) \geq \pi(x, y)$ for every $x, y \in X$ makes consumers weakly better off, because the probability they will choose the cheapest alternative can only increase.

Does this simple intuition extend to equilibrium analysis? When (X, π') is weighted-regular, the answer is affirmative. As we saw, under weighted regularity both firms earn max-min payoffs. Clearly, greater comparability lowers the max-min payoff, because it raises the seeker's equilibrium payoff in the hide-and-seek game.

The answer is different when (X, π') is not weighted-regular. Consider the case of bi-symmetric CSs that violate weighted regularity, where equilibrium payoffs are given

[8] To check that the strategy given by (9)-(11) is indeed a symmetric equilibrium strategy, all we need to do is verify that firms weakly prefer adopting formats in Z (Y) conditional on charging $p \leq p^m$ $(p \geq p^m)$. We leave this task to the reader.

by (12). Imagine a regulator who wishes to impose a product description standard that will enhance comparability. Suppose that $q_W^* < q < q_Z^*$. If the regulator's intervention increases the values of q or q_W^*, the intervention will indeed lower equilibrium profits. If, however, the intervention causes an increase in the value of q_Z^* without changing q and q_W^* - for instance, by merging all formats in Z into a single, "harmonized" format - the intervention will *raise* equilibrium profits (without affecting the max-min payoff).

The intuition is as follows. In the cutoff equilibrium, the probability that a firm charging $p = 1$ faces a price comparison is a weighted average of q and q_W^*. The parameter q_Z^* affects this probability only indirectly, by changing the equilibrium weights. Specifically, a higher q_Z^* gives expensive firms a stronger incentive to adopt the "complex" formats that constitute W. As a result, the equilibrium cutoff price p^m changes and firms are more likely to charge a price above p^m and thus adopt the W formats. Since the intervention leaves q and q_W^* unchanged, and since $q > q_W^*$, the overall probability that an expensive firm faces a price comparison decreases. As a result, expensive firms enjoy greater market power de facto. Hence, "local" improvements in comparability may have a detrimental impact on consumer welfare.

6 Consumer Switching

In this section we explore another intuition, captured by the quotes from the Introduction, namely that greater comparability leads to more frequent consumer switching. Consider an arbitrary order-independent CS. In a symmetric equilibrium, the probability with which the consumer switches between firms *conditional on making a price comparison* (a quantity known in the marketing literature as the "*conversion rate*") is $\frac{1}{2}$. The reason is that, by the symmetry of π, the posterior probability distribution over price profiles (p_1, p_2) conditional on making a comparison is symmetric. Therefore, the probability that the consumer's default firm is the more expensive option is $\frac{1}{2}$.

Since the conversion rate is $\frac{1}{2}$, it follows that the switching rate is half the probability that consumers make a price comparison. Under weighted regularity, we saw that the comparison probability is v^*, independently of the prices that firms charge, and therefore the switching rate is $\frac{1}{2}v^*$. Thus, when we compare two weighted-regular CSs, any improvement in comparability leads to a higher switching rate (and, as we saw, lower equilibrium profits). This corroborates the intuition that more frequent switching is associated with greater competitiveness.

When weighted regularity is violated, the situation is different. Consider the case of bi-symmetric CSs. The ex-ante comparison probability in the cutoff equilibrium is

$$[F(p^m)]^2 q_Z^* + 2F(p^m)(1 - F(p^m))q + [1 - F(p^m)]^2 q_W^*$$

The co-movement of this expression with the competitiveness of the market outcome is ambiguous because, as we already showed, equilibrium profits in the relevant parameter range decrease with q_W^* and increase with q_Z^*. Thus, when the firms' equilibrium price and format decisions are correlated, the positive link between the switching rate and market competitiveness may break down.

7 Order-Dependent Comparability Structures

In this section, we relax order independence. In Section 2, we noted that the assumption of order independence makes sense when formats represent measurement units. In other cases, ease of comparison may depend on the order in which the consumer considers the market alternatives. For example, suppose that a format is a partition of some set of "sun spots" (i.e., utility-irrelevant states of Nature), and that each firm presents its price as a function of the partition it adopts. The actual price is obtained by taking an average over all sun spots. Suppose that the partition y is a coarsening of the partition x. The consumer may find comparison easier when his default firm adopts x while the other firm adopts y, than the other way around. In this case, the consumer's final choice probabilities are likely to depend on his initial default assignment.

We begin by extending the notion of weighted regularity to order-dependent CSs.

Definition 2 *A comparability structure (X, π) is weighted-regular if there exist $\beta \in \Delta(X)$ and $\bar{v} \in [0,1]$ such that for all $x \in X$:*

$$\sum_{y \in X} \beta(y)\, \pi(x, y) \geq \bar{v} \geq \sum_{y \in X} \beta(y)\, \pi(y, x)$$

We say that such β verifies weighted regularity.

This definition is equivalent to Definition 1 under order independence. When β verifies weighted regularity, the two inequalities in Definition 2 are binding for every x in the support of β. To see why, suppose, without loss of generality, that the L.H.S inequality in Definition 2 is strict for some $x \in X$ for which $\beta(x) > 0$. Then, we obtain

$$\sum_{x \in X} \sum_{y \in X} \beta(x) \beta(y) \pi(x, y) = v(\beta, \beta) > \bar{v}$$

Now consider the R.H.S inequality in Definition 2. We obtain

$$\sum_{x \in X} \sum_{y \in X} \beta(x) \beta(y) \pi(y, x) = v(\beta, \beta) \leq \bar{v}$$

a contradiction.

Building on this observation, it is possible to establish an equivalent definition of weighted regularity in terms of the associated hide-and-seek game, as in Section 3.1.

Lemma 2 *The distribution $\lambda \in \Delta(X)$ verifies weighted regularity in (X, π) if and only if (λ, λ) is a Nash equilibrium in the associated hide-and-seek game.*

The proof is omitted because it proceeds as the proof of the analogous Lemma 1. As in Section 3.1, this equivalence implies that when β verifies weighted regularity, $\bar{v}$ is equal to v^*, the value of the associated hide-and-seek game.

The link between weighted regularity and max-min equilibrium payoffs, established for order-independent CSs, survives the present extension only in one direction.

Proposition 4 *Suppose that (X, π) satisfies weighted regularity. Then, firms earn max-min payoffs in any symmetric Nash equilibrium.*

The proof follows the same line of reasoning as Theorem 1. Fix a symmetric Nash equilibrium. By weighted regularity, a firm can adopt a format strategy $\beta \in \Delta(X)$ such that: (i) the probability that a consumer who is initially assigned to the firm will make a price comparison is weakly below v^*; and *(ii)* the probability that a consumer who is initially assigned to the opponent will make a price comparison is weakly above v^*. It follows that the firm's market share is bounded from below by the R.H.S of (2), and the proof proceeds exactly as in the case of order independence.

The converse to this result does not hold in general. When an order-independent CS violates weighted regularity, it does not follow that firms necessarily earn payoffs above the max-min level in symmetric equilibrium. For example, let $X = \{a, b\}$, $\pi(a, b) = q$ and $\pi(b, a) = 0$, $0 < q < 1$. This CS violates weighted regularity. However, it admits a symmetric Nash equilibrium in which firms play a format strategy that satisfies $\lambda(a) = \frac{1-q}{2-q}$, and a pricing strategy for which the supports of F^a and F^b are $[\frac{1}{3+q}, 1]$ and $[\frac{1-q}{3-q^2}, \frac{1}{3+q}]$. The marginal format strategy is a max-min strategy for the seeker in the associated hide-and-seek game, and therefore firms earn max-min payoffs.

8 Concluding Remarks

This paper studied the implications of limited comparability for market competition. We constructed a model of price competition in which the probability that a consumer makes a price comparison is a function of the formats that firms employ to describe the price of their product. Our key insight is that a notion of uniform comparability, captured by the formal definition of weighted regularity, is crucial for equilibrium analysis. Several key questions concerning equilibrium behavior - whether firms earn profits in excess of the max-min (i.e., "constrained competitive") level, whether price and format decisions are correlated, and how industry profits and consumer switching react to regulatory interventions - hinge on this property.

We devote the final section to a discussion of several aspects of the model.

Simultaneity of price and format decisions

Our model assumes that firms choose prices and formats simultaneously. One could argue that in many situations, formats tend to be a more permanent fixture than prices. An alternative modeling strategy that addresses this criticism would be to assume that firms compete in prices only after committing to the format. We opted for simultaneity because we believe that in many situations of interest - especially in modern online environments - determining a product's price and how to present it are naturally joint strategic decisions.

At any rate, analyzing the alternative, two-stage model is rather straightforward. For simplicity, consider the case of order-independent CSs. For a given profile (x_1, x_2) of the firms' first-stage format decisions, there is a unique equilibrium in the second-stage subgame, in which both firms mix over prices according to (6), except that v^* is replaced with $\pi(x_1, x_2)$. Each firm earns a payoff of $\frac{1}{2}[1 - \pi(x_1, x_2)]$.

In the first stage, firms make their format decisions as if they play a common-interest game in which they share the payoff function $\frac{1}{2}[1 - \pi(x_1, x_2)]$. In equilibrium, each firm i chooses a format strategy λ_i that minimizes $v(\cdot, \lambda_j)$. It follows that when the CS is weighted-regular, there is always a subgame-perfect equilibrium in which each firm plays in the first stage a format strategy that verifies weighted regularity. Firms earn a payoff of $\frac{1}{2}(1-v^*)$ in such an equilibrium. In this limited sense, our equilibrium analysis for weighted-regular CSs extends to the two-stage variant on our model.

However, the two-stage model can generate additional equilibria. For example, consider the CS represented by Figure 2, and let $q = 1$. This CS is weighted-regular. In our model, it satisfies the condition for a competitive equilibrium outcome. However, in the two-stage model, there exists a symmetric sub-game perfect equilibrium in which firms mix uniformly over all four peripheral formats in the first stage. The two-stage model can also generate asymmetric equilibria for CSs that induce only symmetric equilibria in our model. For example, consider the two-format example in which $X = \{a, b\}$, $\pi(a, b) = \pi(b, a) = 0$. In this case, there is an asymmetric equilibrium in which the two firms choose different formats in the first stage and charge $p = 1$ in the second stage, thus earning a payoff of $\frac{1}{2}$ each.

Exogeneity of the CS

Our model takes the CS as given: the function π represents an exogenous distribution over an unobservable characteristic of the population from which our consumer is drawn, namely the ability to compare formats. We view this as a primitive of the consumer's choice procedure, analogous to preferences.

One could argue that π could be derived as part of an equilibrium in a larger model, in which the consumer (optimally) chooses in a prior stage whether to acquire this ability by incurring "thinking costs". For example, when formats represent measurement units, the consumer's limited ability to convert units could be derived from an earlier decision not to memorize the conversion rates.

Nevertheless, for many purposes it makes sense to regard π as exogenous. Even if the consumer's mastery of conversion rates is a consequence of prior optimization, it is probably obtained taking into account a multitude of market situations, in addition to the one in question. In other words, it is optimization in a "general equilibrium" sense, whereas we focus on a "partial equilibrium" analysis. However, we do believe that endogenizing the CS in this manner is an interesting direction for future work. In particular, it raises the question of whether weighted regularity could emerge as a necessary property of equilibrium in the larger model.

Another dimension of the consumer's choice procedure which could be endogenized is the initial assignment of consumers to firms.[9] One could argue that in the long

[9] We thank an editor of this journal for suggesting this extension.

run, consumers would be inclined to leave firms that adopt complex, opaque formats (and thus increase "thinking costs", to use the terminology of the larger model above). This consideration can be incorporated by requiring that in equilibrium, not only the comparability of formats but also the initial default assignment maximize consumer utility. This added requirement may generate incentives toward greater transparency of formats.

Asymmetric Nash equilibria

Our analysis has focused mainly on symmetric Nash equilibria. We wish to remark that Proposition 1 holds for any equilibria. The proof is similar to the proof for symmetric equilibria, with additional, tedious details due to potential discontinuities. Piccione and Spiegler (2009) prove that the symmetric equilibrium we constructed for the example of m non-comparable formats at the beginning of Section 3 is the only Nash equilibrium in the game.

Equilibria that display asymmetry in the format dimension are easy to generate in many cases. For example, suppose that λ and λ' are both format strategies that verify weighted regularity for a given CS. Then, there exist equilibria in which one firm plays λ while the other plays λ', and both firms mix over prices according to (6), independently of their format choice. We conjecture, however, that for any CS, firms' marginal pricing strategies in any Nash equilibrium are identical. We also conjecture that any CS gives rise to unique equilibrium payoffs. Proving these conjectures is left for future work.

Revealed-preference properties of the consumer's choice procedure

From a procedural point of view, the consumer in our model departs from rational decision making because his behavioral rule contains an element of limited comparability. However, sometimes a choice model can be procedurally non-rational, yet consistent with rational choice from a revealed-preference point of view (a prime example is satisficing - see Rubinstein (1998, p. 24)). To explore the basic revealed-preference properties of our model of consumer choice, restrict attention to deterministic CSs. Define the revealed strict preference relation as follows. The format-price pair (x', p') is revealed to be strictly preferred to the pair (x, p) if and only if $\pi(x, x') = 1$ and $p' < p$ (that is, the consumer switches away from (x, p) to (x', p') when the former is his default). This is a conventional definition in choice models with a default option (see Masatlioglu and Nakajima (2010) or Eliaz and Spiegler (2011)).

As usual, the revealed indifference relation is derived from this relation: the consumer reveals that he is indifferent between (x', p') and (x'', p') if neither pair is revealed to be strictly preferred to the other. Clearly, this relation is reflexive: the consumer never switches away from his default if the other firm plays an identical format-price pair. Obviously, the revealed indifference relation does not represent true indifference, because it typically captures inability to make a comparison.

It is easy to show that consumer choice is consistent with maximizing a utility function over $X \times [0,1]$ if and only if $\pi(x,y) = 1$ for all $x, y \in X$ - in which case, our model collapses to Bertrand competition. Even when we focus attention on the revealed *strict* preference relation, we see that it is typically intransitive. To see why, consider the following order-independent CS: $X = \{a,b,c\}$, $\pi(x,y) = 1$ for all $x, y \in X$ except for $\pi(a,c) = 0$. Suppose that $p < p' < p''$. The pair (a,p) is strictly preferred to the pair (b,p'), which in turn is strictly preferred to (c,p''). Yet, the consumer is revealed to be indifferent between (a,p) and (c,p'').

In fact, it is easy to show that the revealed strict preference relation is transitive if and only if the graph of the CS represents an equivalence relation over X - that is, if $<\pi(x,x') = 1$ and $\pi(x',x'') = 1>$ always implies $\pi(x,x'') = 1$. Symmetric equilibria for such CSs are basically the same as in the example analyzed at the beginning of section 3. Suppose that the equivalence relation has m equivalence classes. Then, the marginal format strategy assigns probability $\frac{1}{m}$ to each equivalence class, and the marginal pricing strategy is given by (1). The only multiplicity that can arise is due to the arbitrariness in assigning probabilities to formats *within* each equivalence class. This observation shows how we can derive a strong equilibrium characterization from a basic restriction on the revealed preference relation induced by our choice procedure.

Taste-driven versus comparability-driven product differentiation

At the end of the Introduction, we noted that formats can be interpreted more broadly than just pricing structures, to encompass other product features that affect comparability. However, this requires a reevaluation of the utility-irrelevance of formats. For example, the package of a product can be viewed as a format that affects comparability, because it can change the probability that consumers notice the product and thus consider it as a potential substitute for their default. At the same time, the package is not necessarily utility-irrelevant, because consumers may derive direct utility from certain aspects of the design.

We are thus led to the question of how our limited-comparability approach is related to conventional models of product differentiation (e.g., see Anderson, de Palma and Thisse (1992)). The firms' mixing over formats in Nash equilibrium of our model can be viewed as a type of product differentiation. Since in our model the firms' product is inherently homogenous (the consumer's willingness to pay is independent of x), such differentiation in formats is purely a reflection of the firms' attempt to avoid price comparisons, whereas in conventional models product differentiation is viewed as the market's response to consumers' differentiated tastes.

To facilitate the comparison between the two approaches, it may be useful to think of our model in spatial terms. Suppose that firms are stores and nodes in the graph representation of a CS represent possible physical locations of stores. A link from one location x to another location y indicates that it is costless to travel from x to y. The absence of a link from x to y means that it is impossible to travel in that direction. According to this interpretation, the consumer follows a myopic search process in which he first goes randomly to one of the two stores (independently of their locations). Then, he travels to the second store if and only if the trip is costless. Finally, the consumer chooses the cheaper firm that his search process has elicited (with a tie-breaking rule that favors the initial firm).

This spatial description is not given here for its realism, but because it is reminiscent of conventional models of spatial competition. However, there is a crucial difference. In conventional models of spatial competition, consumers are attached to specific locations and choose between stores according to their price and the cost of travelling to their location. (To make the analogy as precise as possible, assume that the cost of traveling from x to y is zero if there is a link from x to y, and infinity in the absence of such a link.) Thus, a consumer who is attached to a location x does not care at all about the cost of transportation between two stores if none of them is located at x. In contrast, consumer choice in our model is always sensitive to the probability of a link between the firms' locations. Another way of stating the difference is that in our model consumer choice is typically impossible to rationalize with a random utility function over format-price pairs - see the previous comment regarding the revealed-preference properties of our model. In contrast, conventional models of spatial competition are by construction consistent with a random utility function over price-location pairs. Therefore, consumer choice probabilities in our model are generally inconsistent with a conventional spatial model of consumer choice.

The two models also differ at the level of equilibrium predictions. In particular, recall the anomalous comparative statics of equilibrium profits with respect to link probabilities in bi-symmetric CSs, observed in Section 5. It can be shown that this effect cannot be reproduced in the conventional spatial-competition analogue: for any initial assignment of consumers to the two sets of nodes W and Z (assuming uniform assignment within each set), increasing the connectivity of nodes in the set Z necessarily lowers equilibrium profits. Exploring further the comparison and interaction between taste-driven and comparability-driven differentiation is an interesting direction which we hope to pursue in future research.

References

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

[2] Anderson, S., A. de Palma and J.F. Thisse (1992): *Discrete Choice Theory of Product Differentiation*, MIT Press.

[3] Banks, J. and Z. Oldfield (2007): "Understanding pensions: Cognitive Function, Numerical Ability and Retirement Saving," *Fiscal Studies* **28**, 143-170.

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

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

[6] Chioveanu, I. and J. Zhou (2010): "Price Competition with Strategic Obfuscation", UCL-ELSE Working Paper no. 339/2009.

[7] Choi, J., D. Laibson and B. Madrian (2009): "Reducing the Complexity Costs of 401(k) Participation through Quick Enrollment," in *Developments in the Economics of Aging*, David Wise, editor, pp. 57-88.

[8] EC (2009): "On the Follow up in Retail Financial Services to the Consumer Market Scoreboard," Commission Staff Working Document, http://ec.europa.eu/consumers/rights/docs/swd_retail_fin_services_en.pdf.

[9] Ellison, G. and A. Wolitzky (2008): "A Search Cost Model of Obfuscation", mimeo. MIT.

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

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

[12] Iyengar, S. and M. Lepper (2000): "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology* **79**, 995-1006.

[13] OFT (2008): "Personal Current Accounts in the UK", an OFT market study, http://www.oft.gov.uk/shared_oft/reports/financial_products/OFT1005.pdf.

[14] Masatlioglu, Y. and D. Nakajima (2010): "Choice by Iterative Search", mimeo, University of Michigan.

[15] Piccione, M. and R. Spiegler (2009): "Framing Competition," UCL-ELSE Working Paper no. 336/2009.

[16] Reny P. (1999), On the Existence of Pure and Mixed Strategy Equilibria in Discontinuous Games, *Econometrica* **67**, 1029-1056.

[17] Rubinstein, A. (1998): *Modeling Bounded Rationality*, MIT Press.

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

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

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

[21] Varian, H. (1980): "A Model of Sales," *American Economic Review* **70**, 651-59.

[22] Wilson, C. (2010): "Ordered Search and Equilibrium Obfuscation," *International Journal of Industrial Organization* **28**, 496-506.

Appendix

Consider a symmetric Nash equilibrium in which the equilibrium strategy is $(\lambda, (F^x)_{x\in X})$. Let S^x denote the support of F^x, and let p^{xu} and p^{xl} denote the infimum and supremum of S^x, respectively. Recall that for every $x \in Supp\,(\lambda)$, F^x is continuous. Therefore, $p^{xl} < p^{xu}$ for any $x \in Supp\,(\lambda)$.

Proof of Proposition 1

Define $X^A = \{x \in X : \pi(y,x) = 1 \text{ for all } y \in X\}$. This is the set of "universally comparable formats".

We first show that if firms earn zero profits in equilibrium, X^A must be non-empty. If $\pi(y,x) < 1$ for some $x \in Supp(\lambda)$ and some $y \in X$, then a firm can make positive profits charging $p = 1$ and choosing y. It follows that if firms earn zero profits, then $Supp(\lambda) \subseteq X^A$, hence X^A is non-empty.

We now show that non-emptiness of X^A implies zero equilibrium profits. Suppose that X^A is non-empty and profits are positive. In order to obtain a contradiction, define the following (finite) subset of prices:

$$\begin{aligned}\Theta = \ & \{p : \text{there exists } y \in Supp(\lambda) \text{ such that } p^{yl} = p \text{ and} \\ & \pi(x,y) < 1 \text{ for some } x \in Supp(\lambda) \text{ for which } p^{xu} > p\}\end{aligned}$$

We first show that, since profits are positive, Θ is not empty. Let Z be the set of formats z in $Supp(\lambda)$ for which $p^{zu} = 1$. If $\pi(z,y) = 1$ whenever $z \in Z$ and $y \in Supp(\lambda)$,then profits are obviously equal to zero. Thus, $\pi(z,y) < 1$ for some $z \in Z$ and some $y \in Supp(\lambda)$. Because $p^{yl} < 1 = p^{zu}$, p^{yl} is in Θ.

Define p^* as the lowest price in Θ, and let $y^* \in Supp(\lambda)$ satisfy $p^{y^*l} = p^*$. Since (y^*, p^*) is in the support of the equilibrium strategy, a firm's equilibrium profit can be evaluated using the pure strategy (y^*, p^*) and is therefore equal to

$$\frac{p^*}{2}\left(1 + \sum_{x \in X} \lambda(x)\left[(1 - F^x(p^*))\pi(x,y^*) - F^x(p^*)\pi(y^*,x)\right]\right)$$

If $F^x(p^*) > 0$ for some $x \in Supp(\lambda)$, then $p^{xl} < p^* < p^{y^*u}$. If $\pi(y^*,x) < 1$, p^{xl} belongs to Θ, contradicting the definition of p^*. Therefore, $\pi(y^*,x) = 1$. It follows that

$$F^x(p^*)\pi(y^*,x) = F^x(p^*)$$

for every $x \in Supp(\lambda)$. Let $x^* \in X^A$. If a firm chooses the pure strategy (x^*, p^*), it earns

$$\frac{p^*}{2}\left(1 + \sum_{x \in X} \lambda(x)\left[(1 - F^x(p^*)) - F^x(p^*)\pi(x^*,x)\right]\right)$$

Clearly, $F^x(p^*)\pi(x^*,x) \leq F^x(p^*)$ for every x. By the definition of Θ, $\pi(x,y^*) < 1$ for some $x \in Supp(\lambda)$ for which $p^* < p^{xu}$. By continuity of F^x, $p^* < p^{xu}$ implies $F^x(p^*) < 1$. It follows that (x^*, p^*) generates a strictly higher payoff than (y^*, p^*), a contradiction.

Proof of Proposition 3

Consider a bi-symmetric CS (X,π). Let $v^x(\lambda)$ be the probability that the consumer makes a price comparison conditional on the event that one firm adopts the format x and the other firm mixes over formats according to λ. That is,

$$v^x(\lambda) = \sum_{y\in X} \lambda(y)\,\pi(x,y) \tag{13}$$

Note that by the definition of bi-symmetry, for every $x,x' \in W$ (similarly, for every $x,x' \in Z$), $v^x(\lambda) \geq v^{x'}(\lambda)$ if and only if $\lambda(x) \geq \lambda(x')$. The profits from adopting format x and charging p can be re-written as

$$p \cdot \left(\sum_{y\in X} \lambda(y) \cdot (1 - F^y(p)) \cdot \pi(x,y) + \frac{1}{2}(1 - v^x(\lambda)) \right),$$

that is, a firms choosing x gets the entire market if x is compared to a format y charging a higher price and half the market if no comparison takes place. The proof relies on a series of lemmas.

Lemma 3 $\lambda(x) = \lambda(x')$ *for any* $x,x' \in W$ *or* $x,x' \in Z$.

Proof. Suppose that $\lambda(w) > \lambda(w')$ for some $w,w' \in W$. The pure strategy (w,p^{wu}) is in the support of the equilibrium strategy, and therefore attains the equilibrium profits. Consider the pure strategy (w',p^{wu}). Since the CS is bi-symmetric, $\pi(w,x) = \pi(w',x)$ for any $x \in X\backslash\{w,w'\}$. Thus, conditional upon facing only formats in $X\backslash\{w,w'\}$, the strategies (w,p^{wu}) and (w',p^{wu}) obtain the same profits. Obviously, $F^w(p^{wu}) = 1$. Hence, the difference between the equilibrium profits and the profits obtained using the pure strategy (w',p^{wu}) is equal to

$$p^{wu} \cdot \left(\lambda(w') \cdot \left(1 - F^{w'}(p^{wu})\right) \cdot (q_W - 1) + \frac{1}{2}\left(v^{w'}(\lambda) - v^w(\lambda)\right) \right)$$

Since $\lambda(w) > \lambda(w')$, we have that $v^w(\lambda) > v^{w'}(\lambda)$. Therefore, the above expression is negative, that is, the payoffs from (w',p^{wu}) are higher than the equilibrium payoffs, a contradiction. An analogous argument for Z establishes the claim. ■

Lemma 4 *For any* $p \in [p^l,1]$, $F^x(p) = F^{x'}(p)$ *whenever* $x,x' \in W$ *or* $x,x' \in Z$.

Proof. Suppose that $F^w(p) > F^{w'}(p)$ for $w,w' \in W$ for some $p \in (p^l,1)$. Since by Lemma 3, $\lambda(w) = \lambda(w')$, we have $v^w(\lambda) = v^{w'}(\lambda)$. Let $p \in Supp(F^w)$. Then, the pure strategy (w,p) is a best-reply and therefore attains the equilibrium payoff. Thus, the difference between the equilibrium profits and the profits obtained using the pure strategy (w',p) is equal to

$$p \cdot \lambda(w) \cdot \left(F^{w}(p) - F^{w'}(p)\right) \cdot (q_W - 1)$$

This expression is negative. A contradiction. ■

Lemma 5 $\lambda(x) > 0$ *for all* $x \in X$.

Proof. Suppose that $\lambda(x) = 0$ for some $x \in W$. By Lemma 3, $\lambda(w) = 0$ for all $w \in W$ and λ is a uniform distribution over Z. Therefore, $v^{z}(\lambda) = q_Z^*$ for every $z \in Z$ and $v^{w}(\lambda) = q$ for every $w \in W$. By assumption, $q_Z^* > q$. Therefore, it is profitable for a firm to deviate to a pure strategy $(w,1)$, $w \in W$.

Now suppose that $\lambda(x) = 0$ for some $x \in Z$. By Lemma 3, $\lambda(z) = 0$ for all $z \in Z$ and λ is a uniform distribution over W. Therefore, $v^{w}(\lambda) = q_W^*$ for every $w \in W$ and $v^{z}(\lambda) = q$ for every $z \in Z$. By assumption, $q_W^* < q$. Therefore, it is profitable for a firm to deviate to a pure strategy (z,p^{l}), $z \in Z$. ■

Define

$$\begin{aligned} A &= \lambda(w) \cdot n_W \cdot (q_W^* - q) \\ B &= \lambda(z) \cdot n_Z \cdot (q_Z^* - q) \end{aligned}$$

Observe that, since the CS violates weighted regularity and, by hypothesis, $q_Z^* \geq q_W^*$, $A < 0$ and $B > 0$.

Lemma 6 *For any* $w \in W$ *and* $z \in Z$, $p^{zu} = p^{wl}$.

Proof. By Lemma 4, all formats $w \in W$ have the same F^{w} and all formats $z \in Z$ have the same F^{z}. In particular, S^{w} is identical for all $w \in W$, and S^{z} is identical for all $z \in Z$. From now on, let w and z be arbitrary formats in W and Z, respectively. Since the support of the marginal pricing strategy F is $[p^{l},1]$ and supports are closed sets by definition, we have that $S^{w} \cup S^{z} = [p^{l},1]$ and $S^{w} \cap S^{z} \neq \varnothing$. Conditional on charging the same price p, the formats w and z must generate the same market share. Therefore, for every $p \in S^w \cap S^z$:

$$\lambda(w) \cdot n_W \cdot q_W^* \cdot (1 - F^w(p)) + \lambda(z) \cdot n_Z \cdot q \cdot (1 - F^z(p)) + \tfrac{1}{2}(1 - v^w(\lambda)) =$$
$$\lambda(w) \cdot n_W \cdot q \cdot (1 - F^w(p)) + \lambda(z) \cdot n_Z \cdot q_Z^* \cdot (1 - F^z(p)) + \tfrac{1}{2}(1 - v^z(\lambda))$$

which simplifies to

$$B \cdot (1 - F^z(p)) - A \cdot (1 - F^w(p)) = \frac{v^z(\lambda) - v^w(\lambda)}{2} \tag{14}$$

We will first show that $S^w \cap S^z$ is a singleton. Assume the contrary, that is, there exist prices $p, p' \in S^w \cap S^z$, $p < p'$. Since the marginal pricing *cdf* F is strictly increasing, either $F^w(p) < F^w(p')$ or $F^z(p) < F^z(p')$. Since $A < 0$ and $B > 0$, if equation (14) is satisfied for p, it is violated for p', a contradiction.

When $S^w \cap S^z$ is a singleton, there are two cases: either

$$p^l = p^{zl} < p^{zu} = p^{wl} < p^{wu} = 1$$

or

$$p^l = p^{wl} < p^{wu} = p^{zl} < p^{zu} = 1.$$

Assume the latter. Then, adopting z is optimal at $p = 1$ and adopting w is optimal at p^l. Therefore, $v^z(\lambda) \leq v^w(\lambda)$. Since the CS is not weighted-regular, $v^z(\lambda) < v^w(\lambda)$. However, as $A < 0$ and $B > 0$, the L.H.S of equation (14) evaluated at $p = p^{wu} = p^{zl}$ is positive, a contradiction. Therefore, $p^{zu} = p^{wl}$. ■

By Lemmas 5 and 6, a symmetric Nash equilibrium must be a cutoff equilibrium in which $\lambda^{[p^m,1]}$ is a uniform distribution over W and $\lambda^{[p^l,p^m]}$ is a uniform distribution over Z, where $p^m = p^{zu} = p^{wl}$. To pin down the format strategy λ, we use the equilibrium condition that firms are indifferent between playing $w \in W$ and $z \in Z$ at the cutoff price p^m. This condition is

$$\lambda(w)\, n_W q - \lambda(z)\, n_Z q_Z^* = \lambda(w)\, n_W q_W^* - \lambda(z)\, n_Z q$$

for arbitrary $w \in W$ and $z \in Z$.