Price Competition under Limited Comparability*

Michele Piccione† and Ran Spiegler‡

First version: May 2009
This version: March 2010

Abstract

This paper studies market competition when firms can influence consumers' ability to compare market alternatives, through their choice of price "formats". We introduce random graphs as a tool for modelling limited comparability of formats. Our main results concern the interaction between firms' equilibrium price and format decisions and its implications for industry profits and consumer switching rates. We show that narrow regulatory interventions that aim to facilitate comparisons may have adverse consequences for consumer welfare. Finally, we argue that our limited-comparability approach provides a new perspective into the phenomenon of product differentiation.

1 Introduction

Standard models of market competition assume that consumers are able to form a ranking of all the alternatives they are aware of. The ranking may reflect informational constraints, but it is complete nonetheless. In reality, consumers do not always carry out all the comparisons that "should" be made. Moreover, whether consumers are able to make comparisons often depends on how alternatives are described, or "framed":

*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 and numerous seminar participants. 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.

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

- Prices and quantities may be defined for different units of measurement that consumers often find difficult to convert to a common standard. For example, the repayment structure of a loan can be defined in terms of various time units. Interest on a bank deposit can be presented in various forms. And nutritional contents of a food product can be specified for various units of weight or volume.

- Price schedules in several industries condition on a large number of diverse 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. Different firms often utilize different, partly overlapping classifications in the presentation of their price schedules, and this complicates the task of comparing them.[1]

Marketers and regulators alike have long recognized the importance of limited comparability as an obstacle to market competition. Nutritional information on food product labels is required to conform to rigid formats which include standardized units of measurement.[2] As to regulation of retail financial services, the following quotes from recent consumer protection reports are representative:

> "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..." (EC (2009), p. 4)

> "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), p. 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

[1] Of course, different classifications partly reflect differences in the cost structure and distribution of consumer preferences that the firms face. However, they have the additional consequence of hindering comparisons, and this may be among the reasons firms adopt them in the first place.

[2] See http://en.wikipedia.org/wiki/Nutrition_facts_label.

> 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 analysis is motivated by the following questions: What are the implications of limited comparability for the competitiveness of the market outcome? How do regulatory interventions aimed at enhancing comparability perform when firms respond strategically to these interventions? What is the relationship between the firms' pricing and framing decisions? How does limited comparability affect the consumer's propensity to switch products?

Our model takes textbook Bertrand competition as a starting point, and complicates it in the direction of limited comparability. Two profit-maximizing firms facing a single consumer produce perfect substitutes at zero cost. They play a simultaneous-move game in which each firm i chooses a price p_i and a pricing structure x_i for its product, referred to as a *format*. The price is the actual payment the consumer makes to the firm when he chooses it, whereas the format is the way in which the price is presented to the consumer. The consumer has a unit demand and 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, say firm 1. We interpret the consumer's initial firm assignment as a default option arising from previous consumption decisions. With probability $\pi(x_1, x_2)$, the consumer makes a price comparison and chooses the rival firm's product if strictly cheaper. Otherwise, he buys from the firm 1. When $\pi(x, y) = 1$ for all formats x, y, comparability is perfect and the model collapses to Bertrand competition. When $\pi(x, y) = \pi(y, x)$ for all formats x, y - a property we dub "*order independence*" - price comparisons are independent of the order in which the consumer considers alternatives.

The consumer's decision procedure exhibits prudence, or "inertia". Whenever the consumer is unable to compare his default option to a new alternative, he chooses the former. Consequently, when the consumer is initially assigned to firm i, he selects it with probability one when $p_j \geq p_i$ and with probability $1 - \pi(x_i, x_j)$ when $p_j < p_i$. This bias in favor of the default is consistent with the notion that, when consumers face complex decision problems, they are likely to fall back on a default option, if they have

one. This behavioral trait has received experimental support (see, for example, Iyengar and Lepper (2000) and Iyengar, Huberman and Jiang (2004)) and appears to be highly realistic in market contexts. In industries such as communication, electricity or retail banking, consumers tend not to switch away from their current (default) provider when comparison is difficult. Indeed, the above-cited consumer protection reports emphasize consumer inertia driven by limited comparability as a major cause of low switching rates and weak competitive forces in these industries.

We represent the comparability structure π as a random graph, where the set of nodes corresponds to the set of formats, and $\pi(x, y)$ is the probability of a directed link from node x to node y. A link from format x to format y means that y is easy to compare to x. The graph representation entails no loss of generality: its role is to visualize comparability structures that involve many formats, suggest fruitful notions of comparability and simplify the exposition of results. By allowing the graph to be probabilistic, we capture heterogeneity among consumers, in that $\pi(x, y)$ can be viewed as the *firms' (common) belief* over the consumer's ability to compare y to x.

It should be emphasized that while our model is partly motivated by practical questions that concern consumer protection, our contribution here is theoretical. Our aim is to obtain an abstract understanding of the implications of limited comparability on market competition and consumer welfare. To attain this objective, we work with the simplest possible market model and sacrifice realism in every dimension in order to attain considerable generality in the novel feature of the model, namely consumers' limited and frame-sensitive ability to compare between alternatives. For this reason, our model cannot be descriptively faithful to any real-life industry.

1.1 An Illustrative Example: The "Star" Graph

We use a simple example to illustrate the model and some of our main insights. Consider a product that can be priced in $m+1$ different currencies, one major and m minor ones. The consumer is able to compare prices denominated in different currencies only if he knows the exchange rate. Let q be the probability that the consumer knows the exchange rate between the major currency and any minor one (whether there is correlation between minor currencies is immaterial). For simplicity, let's assume that the consumer does not know the exchange rates between the minor currencies. The resulting comparability structure can be represented as a "star" graph, such as the one

given by Figure 1:

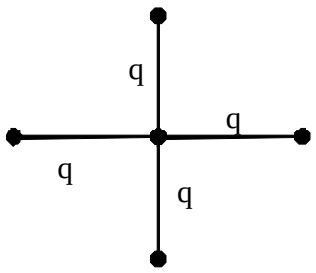

(Figure 1)

A star graph has one "core" node, representing prices denominated in the major currency, and m "peripheral" nodes ($m = 4$ in Figure 1) representing prices denominated in a minor currency. Every node is linked to itself with probability one. In addition, the core node is linked to each of the "peripheral" nodes with probability $q \in (0, 1)$.[3]

The star graph admits no pure-strategy Nash equilibrium. On one hand, a perfectly competitive outcome with zero profits is inconsistent with equilibrium because when a firm charges a price $p > 0$ and randomizes over all peripheral formats, it ensures that, with positive probability, the consumer will fail to make a price comparison. On the other hand, a non-competitive outcome is inconsistent with pure-strategy equilibrium by a simple undercutting argument. Since every format is perfectly comparable to itself, a firm can always mimic its opponent's format and slightly undercut its price. Thus, equilibrium strategies are necessarily mixtures over price-format pairs, reflecting a dispersion of prices and formats in the market. The question is how these two components are related.

Symmetric mixed-strategy Nash equilibrium is unique. Its structure turns out to depend on the expected number of minor currencies the consumer knows how to convert into the major currency (and vice versa). When $mq > 1$, the firms' price and format decisions are perfectly correlated. Specifically, there exists a cutoff price p^m, such that firms adopt the core format with probability one conditional on charging a price below p^m, and firms randomize uniformly over all peripheral formats conditional on charging

[3]In this paper, diagrams that represent order-independent graphs are drawn as non-directed graphs and not as directed graphs with symmetric link probabilities. The difference is that in the latter, the link between x and y is realized independently of the link between y and x, whereas in the former they are realized simultaneously. The two are payoff-equivalent for firms. In addition, throughout the paper, diagrams suppress self-links.

a price above p^m. In contrast, when $mq \leq 1$, the firms' pricing decisions are identical across formats. In particular, the equilibrium marginal (mixed) format strategy λ^* has the property that when one firm plays λ^*, its rival is indifferent among all formats because they all induce the same probability of a price comparison.

The threshold $q^* = \frac{1}{m}$ is of interest. When $mq > 1$, the core format *dominates* peripheral formats in terms of comparability, in that adopting it leads to a higher comparison probability *regardless* of the rival firm's format decision. In contrast, when $mq < 1$, each format can induce a higher probability of a price comparison, depending on the rival firm's format strategy. The equilibrium format strategy λ^* is precisely the distribution that equalizes the probability of price comparison across formats, thus neutralizing the relevance of format decisions for comparability.

The equilibria in the two parameter regions are also fundamentally different in terms of industry profits. When gauging the competitiveness of a market outcome, our benchmark is *max-min profits*: each firm earns the minimal profit enabled by the consumers' bounded rationality as defined by the comparability structure. When $mq > 1$, firms earn equilibrium profits above the max-min level. When firm 1 charges the highest price in the equilibrium distribution (equal to the consumer's reservation value), it adopts a peripheral format because it minimizes comparability. For firm 2 to act as competitively as possible (so as to push firm 1's payoff to the max-min level), it should adopt the core format because it maximizes comparability. In equilibrium, however, whenever firm 2 charges a price above p^m, it adopts the less comparable, peripheral formats, thus lowering the overall probability of price comparison and giving firm 1 additional market power which yields profits in excess of the max-min level. In contrast, when $mq \leq 1$, equilibrium profits achieve the max-min level by a straightforward application of the Minimax Theorem. In particular, the fact that the equilibrium format strategy λ^* induces a comparison probability that is independent of the opponent's format choice implies that λ^* both max-minimizes and min-maximizes the probability of a price comparison. As a result, when a firm charges the reservation value (the highest price in the equilibrium distribution), it earns max-min profits.

What are the theoretical implications of this equilibrium analysis for market regulation? Current regulatory practice seeks to harmonize product description and minimize the number of formats. Indeed, in the case of the star graph, industry profits and expected prices increase with m and decrease with q. This is consistent with the intuition that simplifying comparison is beneficial for consumer welfare. However, as we shall see later, in environments that are only slightly more complex than the star graph, this intuition can be misleading, and regulatory interventions that enhance comparability

can make the market outcome less competitive, once the firms' equilibrium response to the intervention is taken into account. In addition, we will show that there is a subtle, non-trivial connection between comparability and the amount of consumer switching that occurs in equilibrium.

1.2 Overview of the Main Results

After presenting the model in Section 2, we analyze Nash equilibria for general order-independent graphs in Section 3. The analysis highlights a novel graph-theoretic property, called "*weighted regularity*", which extends the familiar notion of regular graphs, and turns out to be the appropriate way to generalize the distinction between the $mq > 1$ and $mq \leq 1$ regions made in the context of the star graph. A graph is weighted-regular if nodes can be assigned weights, such that each node has the same total weighted expected number of links. Under weighted regularity, all formats are equally comparable, once the frequency with which they are used is factored in.

We show that if a graph is weighted-regular, there exists a symmetric Nash equilibrium in which the firms' price and format strategies are statistically independent, and their payoffs are equal to the max-min level. Conversely, if firms' price and format strategies are statistically independent in some Nash equilibrium, the graph must be weighted-regular and firms necessarily earn max-min payoffs in this equilibrium. Moreover, their marginal pricing strategies must be identical. Thus, correlation between price and format decisions is a necessary (observable) manifestation of "collusive" equilibrium profits.

In Section 4, we turn to a class of order-independent graphs, referred to as "*bi-symmetric*", which generalize the star graph. In bi-symmetric graphs, the set of formats is partitioned into two categories, such that the probability of a link between two formats depends only on their categories. We provide a complete characterization of the (unique) symmetric Nash equilibrium for bi-symmetric graphs. We use this characterization to demonstrate that regulatory interventions that enhance comparability may have subtle and unexpected implications for equilibrium profits and consumer switching.

In Section 5, we relax order independence and examine the extent to which our equilibrium characterization for order-independent graphs can be extended. Section 6 is devoted to a discussion of the relation between our model and the more conventional view of product differentiation based on preference heterogeneity.

1.3 Related Literature

Our paper joins recent attempts to formalize in broad terms the various ways in which choice behavior is sensitive to the "framing" of alternatives. Rubinstein and Salant (2008) study choice behavior, where the notion of a choice problem is extended to include both the choice set and a frame, interpreted as observable information which should not affect the rational assessment of alternatives but nonetheless affects choice. A choice function assigns an element in the choice set to every "frame-augmented" choice problem. Rubinstein and Salant conduct a choice-theoretic analysis of such extended choice functions, and identify conditions under which extended choice functions are consistent with utility maximization. Bernheim and Rangel (2007) use a similar framework to extend standard welfare analysis to situations in which choices are sensitive to frames. Our notion of "frame dependence" differs from the one in the above models. First, we associate frames (i.e., formats) with individual alternatives, rather than entire choice sets. Second, in our model framing creates preference incompleteness but never leads to preference reversal. Finally, our focus is on market implications rather than choice-theoretic analysis.

This paper is closely related to Eliaz and Spiegler (2007), which first formalized the idea that framing (and marketing in general) affects preference incompleteness by influencing the set of alternatives that consumers subject to their preference ranking. There are two major differences. First, Eliaz and Spiegler (2007) assume that the consumer's propensity to consider a new market alternative is a function of its frame and the default's payoff-relevant details. Second, in the market applications analyzed in Eliaz and Spiegler (2007), framing decisions are costly and price setting is assumed away. The resulting market model is substantially different from ours, emphasizing the firms' trade-off between increasing their market share and economizing on their fixed marketing costs. Chioveanu and Zhou (2009) analyze a many-firms variant on our model in which the comparability structure is a reduced form of the star graph and consumers lack default options. They show that the market equilibrium need not converge to the competitive outcome as the number of firms tends to infinity.

More generally, our paper contributes to a growing theoretical literature on the market interaction between profit-maximizing firms and boundedly rational consumers. Rubinstein (1993) analyzes monopolistic behavior when consumers differ in their ability to understand complex pricing schedules. Piccione and Rubinstein (2003) study intertemporal pricing when consumers have diverse ability to perceive temporal patterns. Spiegler (2006a,b) analyzes markets in which profit-maximizing firms compete over consumers who rely on naive sampling to evaluate each firm. Gabaix and Laib-

son (2006) and Eliaz and Spiegler (2008) study interaction with consumers having limited awareness of future contingencies. Spiegler (2006b) and Gabaix and Laibson (2006) are specifically preoccupied with ways firms strategically use "confusing" pricing schemes to increase consumers' decision errors. Other papers (Carlin (2008), Ellison and Wolitzky (2008) and Wilson (2008)) model obfuscation as a deliberate attempt to increase rational consumers' search costs.

Finally, our paper can be viewed as an extension of a well-known model due to Varian (1980), in which consumers are divided into two groups: those 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 equilibrium, firms play a mixed pricing strategy. In Varian's model, the fraction of "loyal" consumers is exogenous, whereas in our model it is a function of the formats that firms adopt for their products. An interesting aspect of our analysis is the characterization of cases (captured by the notion of weighted regularity) in which format decisions become irrelevant in equilibrium, such that our model effectively collapses into Varian's.

2 The Model

A *graph* is a pair (X, π), where X is a finite set of nodes and $\pi : X \times X \to [0,1]$ is a function that determines the probability $\pi(x,y)$ with which a directed edge links node x to node y. Let n denote $|X|$. We refer to nodes as *formats*, as they represents various ways in which firms can frame the pricing of an intrinsically homogeneous product. A graph π is *deterministic* if for every distinct $x, y \in X$, $\pi(x,y) \in \{0,1\}$. A graph π is *order independent* if $\pi(x,y) = \pi(y,x)$ for all $x, y \in X$. Assume that $\pi(x,x) = 1$ for every $x \in X$ - that is, every format is linked to itself.[4]

Consider a market consisting of two identical, expected-profit maximizing firms and one consumer. The firms produce a homogenous product at zero cost. The consumer is interested in buying one unit of the product. His willingness to pay for the product is 1, independently of the firms' format decisions. The firms play a simultaneous-move game with complete information. A pure strategy for firm i is a pair (p_i, x_i), where $p_i \in [0,1]$ is a price and $x_i \in X$ is a format. We allow firm i to employ mixed strategies of the form $\left(\lambda_i, (F_i^x)_{x \in Supp(\lambda_i)}\right)$, where $\lambda_i \in \Delta(X)$ and F_i^x is a *cdf* over $[0,1]$ conditional on $x \in Supp(\lambda_i)$. We refer to λ_i as firm i's format strategy and to F_i^x as firm i's pricing strategy at x. The marginal pricing strategy induced by a mixed strategy

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

$\left(\lambda, (F^x)_{x \in Supp(\lambda)}\right)$ is

$$F = \sum_{x \in Supp(\lambda)} \lambda(x) F^x$$

Given a *cdf* F on $[0,1]$, let F^- denote its left limit. For any subset non-empty $Z \subseteq X$, $\mathcal{U}(Z)$ denotes the uniform distribution over Z.

Given a realization $(p_i, x_i)_{i=1,2}$ of the firms' strategies, the consumer chooses a firm according to the following rule. He is randomly assigned to a firm - with probability $\frac{1}{2}$ for each firm. Suppose that he is assigned to firm i. If there is a direct link from x_i to x_j - an event that occurs with probability $\pi(x_i, x_j)$ - the consumer makes a price comparison and chooses firm j if $p_j < p_i$. In all other cases, the consumer chooses the initially assigned firm i.

To illustrate the firms' payoff function, consider the graph given by Figure 2. Let $X = \{x, y\}$, $\pi(x, y) = q$ and $\pi(y, x) = 0$. Suppose that firm 1 adopts the format x while firm 2 adopts the format y. If $p_1 < p_2$, firm 1 earns a payoff of $\frac{1}{2}p_1$ while firm 2 earns $\frac{1}{2}p_2$. If $p_1 > p_2$, firm 1 earns $p_1 \cdot (\frac{1}{2} - \frac{1}{2}q)$ while firm 2 earns $p_2 \cdot (\frac{1}{2} + \frac{1}{2}q)$.

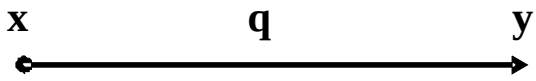

(Figure 2)

When firm i plays the mixed strategy $\left(\lambda_i, (F_i^x)_{x \in Supp(\lambda_i)}\right)$, we can write firm j's expected payoff from the pure strategy (p, x) as follows:

$$\frac{p}{2} \cdot \{1 + \sum_{y \in X} \lambda_i(y) \cdot [(1 - F_i^y(p)) \cdot \pi(y, x) - F_i^{y-}(p) \cdot \pi(x, y)]\}$$

Is consumer choice rational?

Fully rational consumers with perfect ability to make comparisons are represented by a complete graph - i.e. $\pi(x, y) = 1$ for all $x, y \in X$. Rational consumers always make a price comparison, and in this case the model is reduced to standard Bertrand competition. For a typically incomplete graph, the consumer's choice behavior is inconsistent with maximizing a random utility function over price-format pairs.

To see why, consider the following deterministic, order-independent graph: $X = \{a, b, c\}$, $\pi(x, y) = 1$ for all $x, y \in X$ except for $\pi(a, c) = 0$. Suppose that $p < p' < p''$. When faced with the strategy profile $((p, a), (p', b))$, the consumer chooses (p, a) with probability one. Similarly, when faced with the strategy profile $((p', b), (p'', c))$, the consumer chooses (p', b) with probability one. However, when faced with the strategy profile $((p, a), (p'', c))$, the consumer chooses each alternative with probability $\frac{1}{2}$. No random utility function over $[0, 1] \times X$ can rationalize such choice behavior. The reason is that the graph represents a binary relation which is intransitive, and this translates into intransitivity of the implied revealed preference relation over price-format pairs.

In general, our model of consumer choice with deterministic graphs is a special case of incomplete preferences over $[0, 1] \times X$, where both strict and weak preference relations may be intransitive, yet the strict preference relation is acyclic. A probabilistic graph merely represents a distribution over such incomplete preferences.

2.1 Hide and Seek

Our analysis will make use of an auxiliary two-player, zero-sum game, which is a generalization of familiar games such as Matching Pennies. The players (not to be identified with the firms) are referred to as *hider* and *seeker*, denoted h and s. The players 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 $\pi(x_h, x_s)$. We will refer to this game as the *hide-and-seek* game associated with (X, π). Given a mixed-strategy profile (λ_h, λ_s) in this game, the probability that the seeker finds the hider is

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

To see the relevance of this auxiliary game to our model, suppose that firm 1's marginal format and price strategies are λ and F, respectively, where the latter is continuous over the support $[p^l, p^u]$. When firm 2 considers charging the price p^u, it should select a format that minimizes the probability of a price comparison. Hence, it behaves as a hider in the hide-and-seek game, where the seeker's strategy is λ. Similarly, when firm 2 considers charging the price p^l, it should select a format that maximizes the probability of a price comparison. Hence, it behaves as a seeker in the hide-and-seek game, where the hider's strategy is λ. When a firm considers charging an intermediate price, it reasons partly as a hider and partly as a seeker.

The value of the hide-and-seek game is

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

The max-min payoff of a firm in our model is thus $\frac{1}{2}(1 - v^*)$. The reason is that the worst-case scenario for a firm is that its opponent plays $p = 0$ and adopts the seeker's max-min format strategy, to which a best-reply is to play $p = 1$ and minimize the probability of a price comparison.

2.2 Discussion

We devote this sub-section to a discussion of several features of our model.

Irrelevance of prices for comparability
Although our framework is quite general, it does rely on a strong, admittedly problematic assumption: the comparability of market alternatives depends only on their formats, and not on the actual prices. Since the modeler could 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. One may wonder whether we could redefine the set of formats in a way consistent with our assumption. The answer is negative. For example, it is natural to assume that consumers are able to compare two different price pairs if and only if one dominates the other. It can be verified that there does not exist a set of formats which represents this comparability structure and at the same time satisfies the assumption that a firm's format choice does not restrict its set of feasible prices. An interesting generalization would assume that every format $x \in X$ is associated with a set of feasible prices $P(x)$.

Default bias
Although the default bias inherent in the consumer's choice procedure is backed by experimental evidence and everyday intuition, one could contemplate alternative assumptions as to how consumers choose when confronting hard-to-compare formats. For example, they could randomize between firms, or switch away from the default with probability one. It should be emphasized that in the case of order-independent

graphs, these alternative assumptions (as well as any rule that does not discriminate between firms 1 and 2) are *equivalent* for equilibrium analysis, as they all induce the same payoff function for firms; the distinction between them is relevant only for our discussion of consumer switching. Only when order independence is relaxed does the distinction matter for firms' equilibrium behavior.

Reservation values

In our model firms cannot use their format decisions to fool consumers into paying a price above the reservation value, even when they are unable to compare formats. One could argue that if consumers have limited ability to understand the price they are facing, firms should be able to charge prices above their willingness to pay. This difficulty with the interpretation of the reservation value is shared by many market models with boundedly rational consumers. One justification is that there is an implicit ex-post participation constraint, which prevents firms from charging prices above the reservation value. This justification makes particularly strong sense, given the default component of our model. Even if a consumer does not understand the price structure of the default option, he can appreciate whether he actually pays more than his reservation value and quit buying from that firm.

Exogeneity of the comparability structure

Our model takes the comparability structure as given: the function π represents an exogenous distribution over an unobservable characteristic of consumers, namely their ability to compare formats. We view this as a primitive of the consumers' choice procedure, analogous to their preferences. The comparability structure could be derived from a larger decision problem, in which the consumer (optimally) chooses in a prior stage whether to acquire this ability by incurring "thinking costs". For example, in the "star graph" example of Section 1.1, the reason why the consumer fails to convert one currency into another could be his choice not to memorize the exchange rate. However, for many purposes, it makes sense to regard π as exogenous. Even if the consumer's mastery of exchange 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.

Simultaneity of price and format decisions

Our model assumes a firm simultaneously chooses a price and a format. An alternative modeling strategy would be to assume that firms compete in prices only after committing to the format. We opt for the former because we believe that in most situations of

interest - particularly in modern online environments - determining a product's price and how to present it are naturally joint decisions; it would be implausible to allow commitment in formats but not in prices. One could argue that in reality, prices seem to fluctuate more than formats. However, note that real-life price fluctuations are a result of many other factors (such as changing costs), from which our model abstracts.

At any rate, analyzing the alternative, two-sage model is straightforward. For simplicity, consider the case of order-independent graphs. For a given profile (x_1, x_2) of the firms' first-stage format decisions, the price competition second-stage subgame proceeds exactly as in Varian (1980), where the probability that the consumer makes a comparison is fixed at $\pi(x_1, x_2)$. In the first stage, firms make their format decisions as if they play a common-interest game in which both share the payoff function $-\pi$, such that in equilibrium, each firm i chooses a format strategy λ_i that minimizes $v(\cdot, \lambda_j)$. For example, whenever the graph has two formats x and y such that $\pi(x, y) = 0$, it is an equilibrium for one firm to choose x and the other to choose y in the first stage, with both firms playing $p = 1$ in the second stage.

2.3 Basic properties of Nash equilibria

We will conduct a detailed analysis of Nash equilibria in the following sections. In this section, we present two preliminary results. The first characterizes the support of the marginal pricing strategies when both firms make positive profits. The second provides a simple necessary and sufficient condition for the equilibrium outcome to be competitive (that is, both firms charge zero prices).

Proposition 1 *In any Nash equilibrium in which firms make positive profits, there exists a price $p^l \in (0, 1)$ such that, for $i = 1, 2$: (i) the support of F_i is $[p^l, 1]$; (ii) F_i is strictly increasing on $[p^l, 1]$.*

Proposition 2 *Let F_i be a Nash equilibrium marginal pricing strategy for firm $i = 1, 2$. Then, $F_1(0) = F_2(0) = 1$ if and only if there exists a format $x^* \in X$ such that $\pi(x, x^*) = 1$ for every $x \in X$.*

A corollary of Proposition 1 is that if firm i earns the max-min payoff $\frac{1}{2}(1 - v^*)$ in Nash equilibrium, firm j's format strategy conditional on $p < 1$ is a max-min strategy for the seeker in the associated hide-and-seek game.

The proofs of these results rely on price undercutting arguments that are somewhat more subtle than familiar ones. For instance, suppose that firm 1's marginal pricing

strategy has a mass point at some price p^* which belongs to the support of firm 2's marginal pricing strategy. In conventional models of price competition, there is a clear incentive for firm 2 to undercut its price slightly below p^*. In our model, however, price undercutting may have to be accompanied by a change in the format strategy in order to be effective. Adopting a new format strategy may be undesirable for firm 2 because it could change the probability of a price comparison when the realization of firm 1's pricing strategy is $p \neq p^*$.

For the rest of the paper, we assume that the necessary and sufficient condition for a competitive equilibrium outcome is violated.

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

This condition ensures that the firms' max-min payoff is strictly positive - or, equivalently, that the value of the associated hide-and-seek game is strictly below one. Once competitive equilibrium outcomes have been eliminated, any Nash equilibrium must be *mixed*. To see why, assume that each firm i plays a pure strategy (p_i, x_i). If $0 < p_i \leq p_j$, then firm j can profitably deviate to the strategy $(p_i - \varepsilon, x_i)$, where $\varepsilon > 0$ is arbitrarily small. If $p_i = 0$, firm i earns zero profits, contradicting the observation that the firms' max-min payoffs are strictly positive. From now on, we will take it for granted that Nash equilibrium is strictly mixed. We favor a population interpretation of symmetric mixed-strategy equilibrium. According to this interpretation, there is a "sea" of firms and consumers, and the two competing firms that each consumer faces are randomly drawn from this population. A mixed equilibrium strategy represents cross-section price-format dispersion.

3 Nash Equilibrium under Order Independence

In this section, we analyze mixed strategy equilibria for order-independent graphs. The analysis hinges on a notion of "uniform comparability" across formats. From a graph-theoretic point of view, the familiar concept of regular graphs is perhaps the most basic notion of uniform comparability. An order-independent graph is regular if $\sum_{y \in X} \pi(x, y) = \bar{v}$ for all $x \in X$. In a regular graph, all formats are equally comparable in that all formats have the same expected number of links. However, this notion of uniform comparability ignores the frequency with which different formats are adopted. If, for example, x is an isolated node yet both firms adopt it with probability one, the

consumer will make a price comparison with probability one. Hence, a proper notion of uniform comparability should take into account the frequency of adoption of different formats.

Definition 1 *An order-independent graph (X, π) is weighted-regular if there exist $\beta \in \Delta(X)$ and $\bar{v} \in [0,1]$ such that $\sum_{y \in X} \beta(y)\, \pi(x,y) = \bar{v}$ for any $x \in X$. We say that β verifies weighted regularity.*

Regularity thus corresponds to a special case in which the format strategy that verifies weighted regularity is $\mathcal{U}(X)$. Note that the set of distributions that verify weighted regularity is convex. The following are examples of weighted-regular, order-independent graphs.

Example 3.1: Equivalence relations. Consider a deterministic graph that in which $\pi(x,y) = 1$ if and only if $x \sim y$, where $\sim$ is an equivalence relation. Any distribution that assigns equal probability to each equivalence class verifies weighted regularity.

Example 3.2: A cycle with random links. 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 $|y - x| = 1$ or $|y - x| = n - 1$, and $\pi(x,y) = 0$ otherwise. A uniform distribution over all odd-numbered nodes verifies weighted regularity.

Example 3.3: Linear similarity. Consider the following deterministic graph. Let $X = \{1, 2, ..., 3L\}$, where $L \geq 2$ is an integer. For every distinct $x, y \in X$, $\pi(x,y) = 1$ if and only if $|x - y| = 1$. A uniform distribution over the subset $\{3k-1\}_{k=1,...,L}$ verifies weighted regularity.

The star graph of Section 1.1 is weighted regular whenever $mq \leq 1$. Let x_c denote the core node. The format strategy that verifies weighted regularity in this case is λ^*, defined by the following equation, which holds for every peripheral format $x \neq x_c$:

$$\lambda^*(x_c) \cdot 1 + (1 - \lambda^*(x_c)) \cdot q = \lambda^*(x_c) \cdot q + \lambda^*(x) \cdot 1 + (1 - \lambda^*(x_c) - \lambda^*(x)) \cdot 0$$

The L.H.S. is the probability of a price comparison of the format x_c, while the R.H.S. is the probability of a price comparison of any peripheral format $x \neq x_c$.

An equivalent definition of weighted regularity makes use of the auxiliary hide-and-seek game. A graph is weighted-regular if and only if the associated hide-and-seek game has a symmetric Nash equilibrium.

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

Proof. 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. Now 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, very format generates the same probability of a price comparison - namely $\bar{v}$ - against λ. ■

An important feature of our model is that it allows firms to condition their pricing strategy on the format they adopt. It is therefore of interest to know when they choose not to do so in equilibrium, especially as this turns out to have important welfare implications. A mixed strategy $\left(\lambda, (F^x)_{x \in Supp(\lambda)}\right)$ exhibits *price-format independence* if $F^x = F^y$ for any $x, y \in Supp(\lambda)$. The next proposition establishes a link between weighted regularity, price-format independence and equilibrium profits. Define the *cdf*

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

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

Proposition 3 *Consider a graph (X, π).*
(i) Suppose that λ^1 and λ^2 verify weighted regularity. Then, there exists a Nash equilibrium in which firm i, $i = 1, 2$, plays the format strategy λ^i and the pricing strategy $F_i^x \equiv G^*$ for all $x \in X$, and earns max-min payoffs.
(ii) Let $\left(\lambda_i, (F_i^x)_{x \in Supp(\lambda_i)}\right)_{i=1,2}$ be a Nash equilibrium in which both firms' strategies exhibit price-format independence. Then, λ_1 and λ_2 verify weighted regularity, firms earn max-min payoffs, and their marginal pricing strategy is given by 1.

Proof. *(i)* Suppose that firm i plays the format strategy λ^i. By the definition of weighted regularity, every format that the rival firm j may adopt attains the same probability of a price comparison v^* against λ^i. We can thus assume that the probability of a price comparison is exogenously fixed at v^*. Construct a *cdf* F such that every $p \in [\frac{1 - v^*}{1 + v^*}, 1]$ generates the same expected payoff. This leads to the following functional equation:

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

The unique solution is G^*. It is straightforward to verify that no firm would want to deviate to a price $p < \dfrac{1-v^*}{1+v^*}$.

(ii) By assumption, $F_i^x = F_i$ for any $x \in Supp(\lambda_i)$, $i = 1, 2$. Therefore, $x \in \arg\min v(\cdot, \lambda_i)$ for every $x \in Supp(\lambda_j)$ - otherwise, it would be profitable to deviate to the pure strategy $(1, y)$ for some $y \in \arg\min v(\cdot, \lambda_i)$. Similarly, $x \in \arg\max v(\cdot, \lambda_i)$ for every $x \in Supp(\lambda_j)$ - otherwise, it would be profitable to deviate to the pure strategy (p^l, y) for some $y \in \arg\max v(\cdot, \lambda_i)$. It follows that (λ_1, λ_2) and (λ_2, λ_1) are Nash equilibria of the associated hide-and-seek game. Hence, as λ_1 and λ_2 max-minimize as well as min-maximize v, (λ_1, λ_1) and (λ_2, λ_2) are also Nash equilibria of the associated hide-and-seek game. By Lemma 1, both λ_1 and λ_2 verify weighted regularity. Relatively standard arguments (see Proposition 1 in Spiegler (2006)) establish that the equilibrium pricing strategy for each firm must be given by (1) if the probability of a price comparison is exogenously fixed at v^*. ■

Formula (1) is precisely the equilibrium strategy in the two-firm case of Varian's model described in Section 1.3 (Varian (1980)). When firms in our model play a format strategy that verifies weighted regularity, they neutralize the relevance of the format decision because this strategy enforces uniform comparability across formats. Therefore, the model is effectively reduced to Varian's model, which does not incorporate format decisions. It should be noted that, for weighted regularity alone, it suffices that the support of the pricing strategies is the same at all formats in the support.

To demonstrate this result, let us revisit some of the examples presented in the previous sub-section. In Example 3.2, suppose that firm 1 (2) plays a format strategy which is a uniform distribution over all odd-numbered (even-numbered) nodes. Both distributions verify weighted regularity. Suppose further that both firms play independently the pricing strategy given by (1), where $v^* = \frac{2}{n}$. This strategy profile constitutes a Nash equilibrium.

In Example 3.3, suppose that both firms play a format strategy which mixes uniformly over the subset of nodes $\{3k-1\}_{k=1,...,L}$. This distribution verifies weighted regularity. Suppose further that both firms play independently the pricing strategy given by (1), where $v^* = \frac{1}{L}$. This strategy profile constitutes a symmetric Nash equilibrium, in which the consumer makes a price comparison if and only if the firms adopt the same format. In this equilibrium, the formats played with positive probability are "local monopolies": when the consumer faces two different formats, he adheres to his default. Price comparisons occur only when both firms use the same format.

While correlation between prices and formats is necessary for "collusive" equilibrium profits, it is not sufficient. In particular, there exist weighted regular graphs that admit Nash equilibria in which price and format decisions are correlated, and yet firms earn max-min payoffs. This is trivially the case when we take a weighted-regular graph and replicate one of its nodes, such that the new graph contains two distinct formats x, x' with $\pi(x, y) = \pi(x', y)$ for every $y \in X$. In this case, we can construct an equilibrium in which the format strategy verifies weighted regularity (hence firms earn max-min payoffs), yet the format x is associated with low prices while the format x' is associated with high prices. In Piccione and Spiegler (2009) we provide a non-trivial example that does not rely on duplicating nodes.

Unlike the link between price-format independence and weighted regularity, we conjecture that there is a logical equivalence between weighted regularity and max-min equilibrium payoffs. However, at present we can only prove partial results that are consistent with this conjecture. For example, suppose that we impose the restriction that at least one firm plays an equilibrium format strategy that has full support.

Proposition 4 *Consider a Nash equilibrium* $\left(\lambda_i, (F_i^x)_{x \in Supp(\lambda_i)}\right)_{i=1,2}$. *If firm 1 earns max-min payoffs and* $\lambda_2(x) > 0$ *for all* $x \in X$, *then* (X, π) *is weighted-regular.*

The proof of this result relies entirely on the associated hide-and-seek game. It shows that if the seeker in the hide-and-seek game has a max-min strategy with full support, there must exist a symmetric Nash equilibrium in this game.

4 Bi-Symmetric Graphs

In this section, we provide a complete characterization of symmetric Nash equilibria in a special class of graphs, which extends the star graph of Section 1.1. An order-independent graph (X, π) is *bi-symmetric* if X can be partitioned into two sets, Y and Z, such that for every distinct $x, y \in X$:

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

where $\max\{q_Y, q_Z, q\} < 1$. Let n_I denote the number of formats in category $I \in \{Y, Z\}$. In the star graph, $n_Z = 1$, $n_Y = m$, $q_Z = 1$ and $q_Y = 0$.

Bi-symmetric graphs are attractive because with simple parameter restrictions they capture various instances of comparability. When $q < \min\{q_Y, q_Z\}$, we may interpret formats within each of the two categories Y and Z as relatively similar and therefore relatively easy to compare, whereas formats from different categories as more difficult to compare. In contrast, when $q_Y < q < q_Z$, we may interpret the formats in category Z as inherently simpler than those in Y (possibly because they contain translations or conversion guides that are absent from the formats in Y).

Define the "average connectivity" within category $I \in \{Y, Z\}$ as

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

Without loss of generality, assume $q_Z^* \geq q_Y^*$.

One can verify (see the proof of Proposition 5 in the Appendix) that a bi-symmetric graph is weighted-regular if and only if

$$(q_Y^* - q)(q_Z^* - q) \geq 0$$

The star graph satisfies $q_Z^* = 1$ and $q_Y^* = \frac{1}{m}$, such that this inequality holds if and only if $mq \leq 1$. When $q_Y^* = q_Z^* = q$, there is a continuum of format strategies that verify weighted regularity. Otherwise, the unique format strategy that verifies weighted regularity assigns probability $(q_Y^* - q)/[(q_Y^* - q) + (q_Z^* - q)]$ to the set Z, and mixes uniformly within Y and within Z. The value of the hide-and-seek game under weighted regularity is

$$v^* = \begin{cases} q \quad \text{when } q_Y^* = q_Z^* = q \\ \\ \dfrac{q_Y^* q_Z^* - q^2}{(q_Y^* - q) + (q_Z^* - q)} \quad \text{otherwise} \end{cases} \tag{2}$$

By Proposition 3, if weighted regularity holds, there is a symmetric equilibrium in which the firm's marginal format strategy verifies weighted regularity, while their (format-independent) pricing strategy is given by (1).

When the condition for weighted regularity is not satisfied - i.e., when q is strictly between q_Y^* and q_Z^* - the value of the hide-and-seek game is $v^* = q$, since there is a Nash equilibrium in this game in which the seeker (hider) plays $\mathcal{U}(Z)$ ($\mathcal{U}(Y)$). It can be verified that there exists a equilibrium with the following "cutoff" structure. 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 $\lambda_L \equiv \mathcal{U}(Z)$, the format strategy conditional on any price $p \in (p^m, 1]$ is $\lambda_U \equiv \mathcal{U}(Y)$. The marginal pricing strategy F satisfies:

$$F(p^m) = \frac{q - q_Y^*}{q_Z^* - q_Y^*} \tag{3}$$

To compute the firms' equilibrium payoff, let us write down the payoff that a firm earns when it plays the pure strategy - which belongs to the support of the equilibrium mixed strategy - consisting of the price $p = 1$ and some format $y \in Y$:

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

Plugging in (3), we obtain the expression:

$$\frac{1}{2} \cdot [\frac{q - q_Y^*}{q_Z^* - q_Y^*} \cdot (1 - q) + \frac{q_Z^* - q}{q_Z^* - q_Y^*} \cdot (1 - q_Y^*)] \tag{4}$$

which strictly exceeds the max-min payoff $\frac{1}{2}(1 - q)$. We omit the full description of the conditional pricing strategies for the sake of brevity. The following proposition characterizes the symmetric equilibria of bi-symmetric graphs.

Proposition 5 *Let (X, π) be a bi-symmetric graph. In any symmetric Nash equilibrium:*
(i) If $(q_Y^* - q)(q_Z^* - q) \geq 0$, firms play a format strategy that verifies weighted regularity. In particular, if $(q_Y^* - q)(q_Z^* - q) > 0$, the pricing strategy at each $x \in X$ is given by (1), where v^* is given by (2). Firms earn the max-min payoff $\frac{1}{2}(1 - v^*)$.
(ii) If $(q_Y^* - q)(q_Z^* - q) < 0$, firms play the cutoff equilibrium characterized by $(\lambda_L, \lambda_H, F)$ above. Their equilibrium payoff is given by (4).

Thus, when parameter values fit situations in which the categorization of formats captures their relative complexity, the firms' equilibrium strategy displays perfect price-format correlation and firms earn "collusive" profits. In contrast, when parameter values fit situations in which the categorization of formats captures their similarity, the equilibrium strategy displays price-format independence and firms earn max-min payoffs.

4.1 Does Greater Comparability Imply More Competitive Outcomes?

Proposition 5 has interesting implications for relationship between industry profits (equivalently, the expected price paid by the consumer) and comparability. Imagine a regulator who wishes to impose a product description standard that will enhance comparability. Suppose that $q_Y^* < q < q_Z^*$. If the regulator's intervention increases the values of q and q_Y^*, the intervention will lower equilibrium profits. If, however, the intervention causes an increase in the value of q_Z^* (without changing q and q_Y^*), the intervention will *raise* equilibrium profits.

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_Y^*. 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 "hiding" formats that constitute Y. 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 format strategy λ_H). Since the intervention leaves q and q_Y^* unchanged, and since $q > q_Y^*$, the overall probability that an expensive firm faces a price comparison decreases. Hence, expensive firms gain greater market power and greater profits.

Thus, "local" improvements in comparability may have a counter-intuitive, detrimental impact on consumer welfare. Finding a general characterization of the class of transformations of π that lead to unambiguously more competitive outcomes is a challenging comparative-statics problem. For instance, in weighted-regular graphs, equilibrium profits unambiguously decrease with π. The reason is that the equilibrium profits are $\frac{1}{2}(1 - v^*)$, where v^* is the value of the hide-and-seek, which increases whenever any entry in the seeker's payoff function is increased.

4.2 Consumer Switching

The consumer protection reports quoted in the Introduction convey the message that greater market competitiveness goes hand-in-hand with consumers switching more frequently, and that limited comparability plays a significant role in this regard. The case of bi-symmetric graphs illustrates some subtleties in the relationship between comparability and switching.

In a symmetric equilibrium, the probability with which the consumer switches firm *conditional on making a price comparison* (a quantity known in the marketing literature as the "*conversion rate*") is $\frac{1}{2}$. The reason is simple. Conditional on making

a comparison, the consumer faces a symmetric posterior probability distribution over price profiles (p_1, p_2). Since the marginal equilibrium pricing strategy is continuous, the probability that the default 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. When the bi-symmetric graph is weighted-regular, the equilibrium strategy displays price-format independence, and the probability of a price comparison is given by expression (2). Since equilibrium payoffs are equal to the max-min level in this range of parameter values, any improvement in comparability leads to a higher switching rate and lower equilibrium profits.

In contrast, when the bi-symmetric graph is not weighted-regular, the equilibrium probability of price comparison is

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

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 increase with q_Y^* and decrease with q_Z^*. Thus, when prices and formats are correlated, the positive link between the switching rate and market competitiveness breaks down.

5 Order-Dependent Graphs

In this section we explore some properties of Nash equilibria for graphs that violate order independence. We begin by extending the notion of weighted regularity.

Definition 2 *A graph (X, π) is weighted-regular if there exist $\beta \in \Delta(X)$ and $\bar{v} \in [0, 1]$ such that $\sum_{y \in X} \beta(y)\pi(x, y) = \sum_{y \in X} \beta(y)\pi(y, x) = \bar{v}$ for all $x \in X$. We say that β verifies weighted regularity.*

The equivalence between weighted regularity and the existence of symmetric equilibrium in the associated hide-and-seek game, established for order-independent graphs, needs to be qualified when order independence is relaxed.

Lemma 2 *(i) If λ verifies weighted regularity, then (λ, λ) is a Nash equilibrium in the hide-and-seek game; (ii) If (λ, λ) is a Nash equilibrium in the hide-and-seek game and $\lambda(x) > 0$ for every $x \in X$, then λ verifies weighted regularity.*

Proof. The proof of part (i) is identical to the order-independent case. Let us turn to part (ii). Suppose that (λ, λ) is a Nash equilibrium in the hide-and-seek game. Let $\mu^x \in \Delta(X)$ denote a degenerate probability distribution that assigns probability one to node x. Since λ is a best-reply for the hider against λ, $v(\mu^x, \lambda) \geq v(\lambda, \lambda)$ for every $x \in X$. By the full-support assumption, if there is a frame $x \in X$ for which $v(\mu^x, \lambda) > v(\lambda, \lambda)$, then $\sum_{x \in X} \lambda(x) v(\mu^x, \lambda) > v(\lambda, \lambda)$. The L.H.S. of this inequality is by definition $v(\lambda, \lambda)$, a contradiction. Similarly, since λ is a best-reply for the seeker against λ, $v(\lambda, \mu^x) \leq v(\lambda, \lambda)$ for every $x \in X$. By the full-support assumption, if there is a frame $x \in X$ for which $v(\lambda, \mu^x) < v(\lambda, \lambda)$, then $\sum_{x \in X} \lambda(x) v(\lambda, \mu^x) < v(\lambda, \lambda)$, a contradiction. It follows that for every $x \in X$, $v(\mu^x, \lambda) = v(\lambda, \mu^x) = v(\lambda, \lambda)$. ■

To see how the full support assumption is necessary for the second part of this lemma, consider the following example. Let $X = \{a, b, c\}$, $\pi(a, b) = \pi(a, c) = 1$ and $\pi(x, y) = 0$ for all other distinct x, y. The hide-and-seek game induced by this graph has a symmetric Nash equilibrium in which both the hider and the seeker play b and c with probability $\frac{1}{2}$ each. However, the graph is not weighted-regular.

The full-support qualification carries over to the next result, which is a variation on Proposition 3. The proof is close as well, and therefore omitted.

Proposition 6 *(i) Suppose that λ^1 and λ^2 verify weighted regularity. Then, there exists a Nash equilibrium in which each firm $i = 1, 2$ plays the format strategy λ^i and the pricing strategy $F_i^x \equiv G^*$ for all $x \in X$, and earns max-min payoffs.*
(ii) Let $\left(\lambda_i, (F_i^x)_{x \in Supp(\lambda_i)}\right)_{i=1,2}$ be a Nash equilibrium in which both firms' strategies exhibit price-format independence and the format strategies have full support. Then, λ_1 and λ_2 verify weighted regularity, firms earn max-min payoffs, and their marginal pricing strategy is given by (1).

One can extend the notion of bi-symmetric graphs by allowing asymmetric connectivity between the sets Y and Z - that is, $\pi(y, z) = q_{YZ}$ and $\pi(z, y) = q_{ZY}$ for every $y \in Y$, $z \in Z$, where $q_{YZ} \neq q_{ZY}$ (while maintaining the assumption that connectivity is symmetric and constant *within* each of the two sets). The reader can easily verify that such graphs are never weighted regular. It turns out that these graphs can give rise to patterns of price-format correlation that are different from those captured by the cutoff equilibria of Section 4. Recall the graph given by Figure 2: $X = \{x, y\}$, $\pi(x, y) = q$ and $\pi(y, x) = 0$. There is a symmetric Nash equilibrium in which the firms play a format strategy that satisfies $\lambda(x) = \frac{1-q}{2-q}$, and a pricing strategy for which the supports of F^x

and F^y are $[\frac{1}{3+q}, 1]$ and $[\frac{1-q}{3-q^2}, \frac{1}{3+q}]$. Thus, the supports of the format-dependent price strategies are nested in one another. Firms earn max-min payoffs in this equilibrium.

6 Asymmetric Firm Assignment

Equilibrium analysis under order dependence is greatly simplified if the assumption that the consumer's initial firm assignment is symmetric is dropped. Suppose that the consumer is initially assigned to firm 1, referred to as the Incumbent. Firm 2 is referred to as the Entrant. In this case, firm 1's max-min payoff is $1 - v^*$, while firm 2's max-min payoff is zero.

Proposition 7 *Any Nash equilibrium* $\left(\lambda_i, (F_i^x)_{x \in Supp(\lambda_i)}\right)_{i=1,2}$ *of the Incumbent-Entrant model has the following properties:*
(i) (λ_1, λ_2) *constitutes a Nash equilibrium in the associated hide-and-seek game in which firm* 1 *(*2*) is the hider (seeker).*
(ii) Firm 1*'s equilibrium payoff is* $1 - v^*$ *while firm* 2*'s equilibrium payoff is* $v^*(1 - v^*)$.
(iii) The firms' marginal pricing strategies over $[1 - v^*, 1)$ *are given by:*

$$\begin{aligned} F_1(p) &= 1 - \frac{1 - v^*}{p} \\ F_2(p) &= \frac{1}{v^*} \cdot [1 - \frac{1 - v^*}{p}] \end{aligned}$$

and F_1 *has an atom of size* $1 - v^*$ *at* $p = 1$.

The simplicity of the equilibrium characterization in this case results from the firms' unambiguous incentives when choosing their format strategies. The Incumbent has an unequivocal incentive to avoid a price comparison (because then it is chosen with probability one), while the Entrant has an unequivocal incentive to enforce a price comparison (because otherwise it is chosen with probability zero).

7 Concluding Remarks: Framing and Product Differentiation

This paper studies the implications of limited, format-sensitive comparability for market competition. Throughout the paper, we adopted a complexity-based interpretation

of the comparability structure. A format was interpreted as a way of presenting prices, and the function π measured the "ease of comparison" between price formats.

However, building on Eliaz and Spiegler (2007), we can offer a broader interpretation of the graph (X, π) and interpret a format as any utility-irrelevant aspect of the product's presentation which affects the propensity to make a preference comparison. In particular, a format can represent an advertising message, a package design or a positioning strategy. According to this interpretation, a link from x to y can mean that the format x reminds the consumer of the format y, or creates mental associations that eventually lead him to pay attention to any product framed by y. From this point of view, our framework is applicable to many aspects of marketing and framing.

However, adopting this broader interpretation of formats makes the assumption that formats are utility-irrelevant less obvious. For example, while the package of a new product may affect the probability that consumers notice it and thus consider it as a potential substitute for their default product, consumers may also derive direct utility from certain aspects of the package design. We are thus led to a comparison between our limited-comparability approach and 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, such differentiation in formats is a pure reflection of the firms' attempt to avoid price comparisons. By comparison, in conventional models product differentiation is viewed as the market's response to consumers' differentiated tastes.

To understand the relationship between the two approaches, it may be useful to think of our model in spatial terms. Suppose that firms are stores and graph nodes 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 re-interpretation 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. In particular, a consumer who is attached to a location x does not care at all about the cost of transportation between two stores if none are located at x. In contrast, consumer choice in our model is always sensitive to the probability of a link between the firms' locations. Recall that in our model consumer choice is typically impossible to rationalize with a random utility function over pairs (p, x). In contrast, conventional models of spatial competition (and product differentiation in general) are by construction consistent with a random utility function over price-location pairs.

Our model and the more conventional spatial-competition analogue are also different at the level of equilibrium predictions. Consider the star graph with $q = 0$. The conventional model admits asymmetric equilibria in which firms adopt different nodes and charge $p = 1$. In contrast, recall that our model rules out pure-strategy equilibria that sustain non-competitive outcomes. In addition, it can be shown that the anomalous comparative statics of equilibrium profits with respect to link strength in bi-symmetric graphs cannot be reproduced in the conventional spatial-competition analogue of our model.

The two perspectives have very different welfare implications. Consider again the star graph. As the number of peripheral formats m increases, equilibrium profits rise. Thus, increasing the number of formats has an unambiguously negative effect on consumer welfare. In contrast, in a standard differentiated-taste model, increasing the number of available brands has an ambiguous effect. On one hand, it weakens competitive forces and thus raises prices (as in our model). On the hand other, it increases the number of available alternatives and thus raises the maximal utility that each consumer can obtain. This latter feature is absent from the limited-comparability perspective.

The two contrasting approaches to product differentiation can be conveniently integrated. Suppose that a consumer type θ is characterized by two primitives: a graph π_θ and a willingness-to-pay function $u_\theta : X \to \{0, 1\}$. The function u_θ essentially describes the set of product formats (or brands) that type θ likes, whereas the graph π_θ determines the type's ability to make price comparisons. Exploring this model, and particularly its ability to account for real-life consumer behavior data, is an interesting challenge for future work.

References

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

[2] Carlin, B. (2008): "Strategic Price Complexity in Retail Financial Markets," *Journal of Financial Economics*, forthcoming.

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

[4] 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.

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

[6] Eliaz, K. and R. Spiegler (2007): "Consideration Sets and Competitive Marketing," mimeo, UCL and Brown University.

[7] Eliaz, K. and R. Spiegler (2008): "Consumer Optimism and Price Discrimination," *Theoretical Economics* **3**, 459-497.

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

[9] 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.

[10] Iyengar, S., G. Huberman and W. Jiang (2004): "How Much Choice is Too Much: Determinants of Individual Contributions in 401K Retirement Plans," in O.S. Mitchell and S. Utkus (Eds.), *Pension Design and Structure: New Lessons from Behavioral Finance* (pp. 83-95). Oxford: Oxford University Press.

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

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

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

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

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

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

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

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

[19] Wilson, C. (2008): "Ordered Search and Equilibrium Obfuscation," mimeo, University of Oxford.

8 Appendix: Proofs

8.1 Proposition 1

Consider a Nash equilibrium in which firms earn strictly positive payoffs. For each firm $i = 1, 2$, let p_i^l denote the infimum of the support of F_i. Clearly, $p_1^l = p_2^l$. For instance, if $p_1^l < p_2^l$, firm 1 makes higher profits by increasing p_1^l at some node. Hence, let p^l denote the infimum of the support of F_1 and F_2. Since profits are positive, $p^l > 0$. Suppose that there is an interval (p, p'), $p^l < p < p' \leq 1$, such that $F_2(p) = F_2^-(p')$. Without loss of generality, we can assume that $F_2(p'') < F_2(p)$ for $p'' < p$. It follows that $F_1(p) = F_1^-(p')$ since the profits of firm 1 from any strategy (p'', x), $p'' \in (p, p')$, in the support of its equilibrium strategy are strictly lower than the profits from $(p''+\varepsilon, x)$, where $\varepsilon > 0$ is sufficiently small. We now show that there exists no $x \in X$ such that (p, x) is a best-reply for either firm. If neither F_1 nor F_2 have a mass point at p, then firm i can profitably deviate from any $(p - \varepsilon, x)$, where $\varepsilon > 0$ is sufficiently small, to

(p'', x), $p'' \in (p, p')$. Suppose then that F_2^x has a mass point at p for some $x \in X$. Such a mass point is a best-reply for firm 2 only if firm 1 also has a mass point at (p, y) for some y for which $\pi(x, y) > 0$ - otherwise, deviating to $(p + \varepsilon, x)$ would be profitable for firm 2, for a sufficiently small $\varepsilon > 0$. But then firm 1 can profitably deviate from (p, y) to $(p - \varepsilon, y)$ for a sufficiently small $\varepsilon > 0$. This concludes the proof.

8.2 Proposition 2

Define $X^A = \{x \in X : \pi(y, x) = 1 \text{ for all } y \in X\}$. Suppose that $F_1(0) = 1$. Then, firm 1 makes zero profits. It follows that $F_2(0) = 1$ and hence firm 2 also makes zero profits. Obviously, $Supp(\lambda_i) \subseteq X^A$, $i = 1, 2$, as if $\lambda_i(x) > 0$ and $\pi(y, x) < 1$ for some y, firm j can make positive profits charging $p = 1$ and choosing y. Hence, X^A is not empty.

Suppose now that X^A is not empty. If $F_1(0) < 1$, then firm 2 makes positive profits. Thus, $F_2(0) < 1$ and firm 1 also makes positive profits. We first show that it is impossible that $\pi(x, y) = 1$ for all $x \in Supp(\lambda_2)$, $y \in Supp(\lambda_1)$. Assume the contrary. By Proposition 1, the upper bound of the support of F_i is equal to 1 for $i = 1, 2$. Take a node z in the support of λ_2 such that the upper bound of the support of F_i^z is equal to one. The profits of firm 2 are equal to

$$\frac{1}{2} \sum_{x \in X} \left(1 - F_1^{x-}(1)\right) \lambda_1(x)$$

Choosing a price equal to $1 - \varepsilon$ and a node x^* in X^A, firm 2 obtains

$$\frac{(1 - \varepsilon)}{2} \sum_{x \in X} \left(1 - \pi(x^*, x) F_1^x(1 - \varepsilon) + (1 - F_1^x(1 - \varepsilon))\right) \lambda_1(x)$$

Since firm 2's payoff is positive, $F_1^{x-}(1) < 1$ for some $x \in Supp(F_1)$. But then, for ε sufficiently small, the second expression is larger than the first expression, a contradiction.

Now let p^* be the lowest price p in $Supp(F_1) \cup Supp(F_2)$ for which there exist $x \in Supp(\lambda_j)$ and $y \in Supp(\lambda_i)$, where $i \neq j$, such that $p \in Supp(F_i^y)$ and $\pi(x, y) < 1$. Obviously, $p^* > p^l$. Without loss of generality, suppose that $p^* \in Supp(F_2^y)$. Firm 2's payoff from the pure strategy (p^*, y) is

$$\frac{p^*}{2} \sum_{x \in X} \left(1 - \pi(y, x) F_1^{x-}(p^*) + \pi(x, y)(1 - F_1^x(p^*))\right) \lambda_1(x)$$

If firm 2 deviates to the pure strategy $(p^* - \varepsilon, x^*)$, $x^* \in X^A$, it will earn

$$\frac{p^* - \varepsilon}{2} \sum_{x \in X} \left(1 - \pi\left(x^*, x\right) F_1^x\left(p^* - \varepsilon\right) + \left(1 - F_1^x\left(p^* - \varepsilon\right)\right)\right) \lambda_1\left(x\right)$$

By the definition of p^*, if $F_1^{x-}\left(p^*\right) > 0$, then $\pi\left(y, x\right) = 1$. Since $\pi\left(x, y\right) < 1$ for some $x \in Supp\left(\lambda_1\right)$, for ε sufficiently small, the second expression is larger than the first expression, a contradiction.

8.3 Proposition 4

The proof is based on the following version of Farkas' lemma. Let Ω be an $l \times m$ matrix and b an l-dimensional vector. Then, exactly one of the following two statements is true: *(i)* there exists $\beta \in \mathbb{R}^m$ such that $\Omega\beta = b$ and $\beta \geq 0$; *(ii)* there exists $\delta \in \mathbb{R}^l$ such that $\Omega^T\delta \geq 0$ and $b^T\delta < 0$.

Suppose that (X, π) is not weighted-regular. Let us first show that for every $\mu \in \Delta(X)$ such that $\mu\left(x\right) > 0$ for all $x \in X$, there exists $\tilde{\mu} \in \Delta(X)$ such that, for all $y \in X$,

$$\sum_{x \in X} \mu\left(x\right) \pi\left(x, y\right) < \sum_{x \in X} \tilde{\mu}\left(x\right) \pi\left(x, y\right)$$

Order the nodes so that $X = \{1, .., n\}$. Any $\beta \in \Delta(X)$ is thus represented by a row vector $(\beta_1, ..., \beta_n)$. Let Π be a $n \times n$ matrix whose ijth entry is $\pi\left(i, j\right)$. Note that $\Pi = \Pi^T$. Since (X, π) is not weighted-regular, there exist no $\beta \in \mathbb{R}^n$ and $c > 0$ such that $\Pi\beta^T = (c, c, ..., c)^T$. By Farkas' Lemma, there exists a column vector $\delta \in \mathbb{R}^n$ such that $\Pi\delta \geq 0$ and $(c, c, ..., c)\delta < 0$. Since $\pi(i, i) = 1$ for every $i \in \{1, ..., n\}$ and $\pi(i, j) \geq 0$ for all $i, j \in \{1, ..., n\}$, we can modify δ into a column vector $\tilde{\delta}$ such that $\tilde{\delta}_i > \delta_i$ for every i, $\Pi\tilde{\delta} > 0$ and $\sum_i \tilde{\delta}_i = 0$. Let $\mu \in \Delta(X)$ and $\mu(i) > 0$ for every $i \in \{1, ..., n\}$. By the construction of $\tilde{\delta}$, $\tilde{\mu} = \mu + \alpha\tilde{\delta}$ is also a probability distribution over X, for a sufficiently small $\alpha > 0$. Then

$$\Pi\tilde{\mu}^T = \Pi\mu^T + \alpha\Pi\tilde{\delta} > \Pi\mu^T$$

In particular, every component of the vector $\Pi\tilde{\mu}^T$ is strictly larger than the corresponding component of $\Pi\mu^T$.

By hypothesis, $\lambda_2(x) > 0$ for all $x \in X$. We have shown that there exists another format strategy $\tilde{\lambda}$ such that every format $y \in X$ induces a strictly higher probability of a price comparison than λ_2. This contradicts that λ_2 is a max-min strategy.

8.4 Proposition 5

Consider a bi-symmetric graph (X, π). Define

$$\begin{aligned} a &= 1 + q_Y (n_Y - 1) - q n_Y \\ b &= 1 + q_Z (n_Z - 1) - q n_Z \end{aligned}$$

One can verify that weighted regularity holds if and only if the system

$$\begin{bmatrix} a & -b \\ n_Y & n_Z \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

has a non-negative solution - that is, if and only if $ab \geq 0$ (or, equivalently, if and only if $(q_Y^* - q)(q_Z^* - q) \geq 0$).

Let $\left(\lambda, (F^x)_{x \in Supp(\lambda)}\right)$ be a symmetric Nash equilibrium strategy, and let F denote the equilibrium marginal pricing strategy. Let S^x denote the support of F^x, and let p^{xl} and p^{xu} denote the infimum and supremum of S^x. 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, that is,

$$v^x(\lambda) = \sum_{y \in X} \lambda(y)\, \pi(x, y) \tag{5}$$

Note that for every $x, x' \in Y$ (similarly, for every $x, x' \in Z$), $v^x(\lambda) = v^{x'}(\lambda)$ if and only if $\lambda(x) = \lambda(x')$.

The following claims establish Proposition 5.

Lemma 3 *$F(p)$ is continuous on $[p^l, 1]$.*

Proof. It follows from standard arguments, due to the symmetry of equilibrium. ■

Lemma 4 *$\lambda(x) = \lambda(x')$ for any $x, x' \in Y$ or $x, x' \in Z$, $i = 1, 2$.*

Proof. Suppose that $\lambda(x) > \lambda(y)$ for some $x, y \in Y$. Firm i's payoff from the pure strategy (p^{xu}, x) is

$$p^{xu} \left(\begin{array}{c} q_Y \lambda(y) \left(1 - F^y(p^{xu})\right) + \\ \sum_{x \in Y - (x,y)} \left(1 - F^x(p^{xu})\right) q_Y \lambda(x) + \sum_{x \in Z} \left(1 - F^x(p^{xu})\right) q \lambda(x) + \frac{1}{2}\left(1 - v^x(\lambda)\right) \end{array} \right),$$

If the firm deviates to the strategy (p^{xu}, y), it earns

$$p^{xu}\left(\begin{array}{c}\lambda(y)\left(1-F^{y}\left(p^{xu}\right)\right)\\ \sum_{x\in Y-(x,y)}\left(1-F^{x}\left(p^{xu}\right)\right)q_{Y}\lambda(x)+\sum_{x\in Z}\left(1-F^{x}\left(p^{xu}\right)\right)q\lambda(x)+\frac{1}{2}\left(1-v^{y}(\lambda)\right)\end{array}\right).$$

Since $\lambda(x) > \lambda(y)$, $v(\lambda) > v^{y}(\lambda)$, hence the deviation is profitable. An analogous argument for Z establishes the claim. ■

Lemma 5 *For any $p \in [p^{l}, 1]$, $F^{x}(p) = F^{x'}(p)$ whenever $x, x' \in Y$ or $x, x' \in Z$.*

Proof. Suppose that $F^{y}(p) > F^{y'}(p)$ for $y, y' \in Y$. Firm i's payoff from the pure strategy (p, y) is

$$p\left(\begin{array}{c}\left(1-F^{y}(p)\right)\lambda(y)+q_{Y}\left(1-F^{y'}(p)\right)\lambda(y)+\\ \sum_{x\in Y-(y,y')}\left(1-F^{x}(p)\right)q_{Y}\lambda(x)+\sum_{x\in Z}\left(1-F^{x}(p)\right)q\lambda(x)+\frac{1}{2}\left(1-v^{y}(\lambda)\right)\end{array}\right)$$

If the firm deviates to the pure strategy (p, y'), it earns

$$p\left(\begin{array}{c}\left(1-F^{y'}(p)\right)\lambda(y)+q_{Y}\left(1-F^{y}(p)\right)\lambda(y)+\\ \sum_{x\in Y-(y,y')}\left(1-F^{x}(p)\right)q_{Y}\lambda(x)+\sum_{x\in Z}\left(1-F^{x}(p)\right)q\lambda(x)+\frac{1}{2}\left(1-v^{y'}(\lambda)\right)\end{array}\right).$$

By Lemma 4, $\lambda(y) = \lambda(y')$ and therefore $v^{y}(\lambda) = v^{y'}(\lambda)$. It follows that the deviation is profitable. ■

Lemma 6 *If $\lambda(x) = 0$ for some $x \in X$, then λ verifies weighted regularity.*

Proof. Suppose that $\lambda(x) = 0$ for some $x \in Y$. By Lemma 4, λ is a uniform distribution over Z - thus, in particular, $\lambda(y) = 0$ for all $y \in Y$. Therefore, $v^{z}(\lambda) = q_{Z}^{*}$ for every $z \in Z$ and $v^{y}(\lambda) = q$ for every $y \in Y$. If $q_{Z}^{*} \neq q$, it must be profitable to deviate either to the pure strategy $(1, y)$ or to the pure strategy (p^{l}, y). If $q_{Z}^{*} = q$, then λ verifies weighted regularity. ■

Lemma 7 *Suppose that $\lambda(x) > 0$ for all $x \in X$. Then:*
(i) If (X, π) is not weighted-regular, either $p^{yu} = p^{zl}$ or $p^{zu} = p^{yl}$ for any $y \in Y$ and $z \in Z$.
(ii) If $p^{yu} = p^{zl}$ or $p^{zu} = p^{yl}$ for any $y \in Y$ and $z \in Z$, (X, π) is not weighted-regular.

Proof. (i) Suppose that (X,π) is not weighted-regular and $v^z(\lambda) < v^y(\lambda)$. By Lemma 5, the nodes in Y have the same F^y and the nodes in Z have the same F^z. Therefore, $S^y \cap S^z \neq \varnothing$, for any $y \in Y$ and $z \in Z$. The following equations must hold in equilibrium.

$$\lambda(z)\, qn_Z\left(1-F^z\left(p^{yu}\right)\right)+\frac{1}{2}\left(1-v^y(\lambda)\right)=$$
$$\lambda(z)\left(1+q_Z\left(n_Z-1\right)\right)\left(1-F^z\left(p^{yu}\right)\right)+\frac{1}{2}\left(1-v^z(\lambda)\right)$$
$$\lambda(z)\, qn_Z+\left(1+q_Y\left(n_Y-1\right)\right)\lambda(y)\left(\left(1-F^y\left(p^{zl}\right)\right)\right)+\frac{1}{2}\left(1-v^y(\lambda)\right)=$$
$$\lambda(z)\left(1+q_Z\left(n_Z-1\right)\right)+qn_Y\lambda(y)\left(\left(1-F^y\left(p^{zl}\right)\right)\right)+\frac{1}{2}\left(1-v^z(\lambda)\right)$$

which simplify to

$$b\lambda(z)\left(1-F^z\left(p^{yu}\right)\right)=b\lambda(z)-a\lambda(y)\left(1-F^y\left(p^{zl}\right)\right)=\frac{v^z(\lambda)-v^y(\lambda)}{2}$$

Hence, $b<0$. Since the graph is not weighted regular, $a>0$. It can be easily verified that the above equations hold only if $F^z\left(p^{yu}\right)=0$ and $F^y\left(p^{zl}\right)=1$. If $v^z(\lambda)>v^y(\lambda)$, a symmetric argument establishes the claim.

(ii) Suppose that $p^{yu}=p^{zl}$. Note that

$$v^z(\lambda)-v^y(\lambda)=b\lambda(z)-a\lambda(y)$$

In equilibrium

$$b\lambda(z)=\frac{b\lambda(z)-a\lambda(y)}{2}$$

Since $\lambda(y),\lambda(z)>0$, we have $ab<0$. A symmetric argument establishes the claim for the case $p^{zu}=p^{yl}$. ■

Lemma 8 *Suppose that $\lambda(x)>0$ for any $x\in X$. If $p^{yu}\neq p^{zl}$ and $p^{zu}\neq p^{yl}$ for any $y\in Y$ and $z\in Z$, then λ verifies weighted regularity, max-min payoffs are obtained, and $F^z(p)=F^y(p)$ for any $p\in[p^l,1]$.*

Proof. Lemma 7 implies that if $p^{yu}\neq p^{zl}$ and $p^{zu}\neq p^{yl}$ for any $y\in Y$ and $z\in Z$ then the graph is weighted-regular. As in the proof of Lemma 7, the following equilibrium conditions must hold

$$b\lambda(z)\left(1-F^z\left(p^{yu}\right)\right)=\frac{b\lambda(z)-a\lambda(y)}{2}$$

$$b\lambda(z) - a\lambda(y)\left(1 - F^y\left(p^{zl}\right)\right) = \frac{b\lambda(z) - a\lambda(y)}{2}$$

First note that if either $b = 0$ or $a = 0$, then either $\lambda(y) = 0$ or $\lambda(z) = 0$. Hence, suppose that $ab > 0$. Setting $(1 - F^z(p^{yu})) = \sigma$ and $\left(1 - F^y\left(p^{zl}\right)\right) = \delta$, rewrite the system in matrix notation as

$$\begin{bmatrix} b\sigma - \frac{b}{2} & \frac{a}{2} \\ \frac{b}{2} & -a\delta + \frac{a}{2} \end{bmatrix} \begin{bmatrix} \lambda(z) \\ \lambda(y) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

This system has a non-null solution if and only if

$$-\sigma - \delta + 2\sigma\delta + 1 = 0$$

which is only possible, for $0 \leq \delta, \sigma \leq 1$, when $\delta = 1, \sigma = 0$ or $\delta = 0, \sigma = 1$. In the former case, $v_i^z(\lambda) = v_i^y(\lambda)$ and thus λ verifies weighted regularity. In the latter case,

$$b\lambda(z) = \frac{b\lambda(z) - a\lambda(y)}{2}$$

and hence positive solutions for $\lambda(z), \lambda(y)$ do not exist when $ab > 0$. Thus in equilibrium, $F^z(p^{yu}) = 1$, $F^y\left(p^{zl}\right) = 0$, and $v^z(\lambda) = v^y(\lambda)$. Consequently, for any $p \in [p^l, 1]$

$$b\lambda(z)(1 - F^z(p)) = a\lambda(y)(1 - F^y(p))$$

Since $v^z(\lambda) - v^y(\lambda) = b\lambda(z) - a\lambda(y) = 0$, we have $F^z(p) = F^y(p)$. ■

Part (i) of the proposition follows from Lemmas 6, 7, and 8. If $q_Y^* < q < q_Z^*$, then a symmetric Nash equilibrium must be a cutoff equilibrium by Lemmas 6 and 7. Moreover, by Lemma 5, it suffices to consider two cases: either λ_U is a uniform distribution over Y and λ_L is a uniform distribution over Z, or λ_U is a uniform distribution over Z and λ_L is a uniform distribution over Y. To pin down the format strategy λ, we use the equilibrium condition that firms are indifferent between playing $y \in Y$ and $z \in Z$ at the cutoff price p^m ($p^m = p^{zu} = p^{yl}$ in the former case, and $p^m = p^{zl} = p^{yu}$ in the latter case).

In the former case, the condition is given by the equation

$$\lambda(y) n_Y q - \lambda(z) n_Z q_Z^* = \lambda(y) n_Y q_Y^* - \lambda(z) n_Z q$$

for arbitrary $y \in Y$ and $z \in Z$. In the latter case, the condition is given by the equation

$$\lambda(z)\, n_Z q - \lambda(y)\, n_Y q_Y^* = \lambda(z)\, n_Z q_Z^* - \lambda(y)\, n_Y q$$

for arbitrary $y \in Y$ and $z \in Z$. Since $q_Y^* < q < q_Z^*$, the latter case is ruled out, and the former equation yields λ.

8.5 Proposition 7

(i) Whenever $p_1 \leq p_2$, the consumer chooses firm 1 with probability one. Whenever $p_1 > p_2$, the consumer chooses firm 2 if and only if he makes a price comparison. Therefore, for every price p that lies strictly above the infimum of $Supp(F_2)$, firm 1's optimal format minimizes $v(\cdot, \lambda_2^L(p))$, where $\lambda_2^L(p)$ denotes firm 2's format strategy conditional on $p' < p$. Similarly, for every price p that lies strictly below the supremum of $Supp(F_1)$, firm 2's optimal format maximizes $v(\lambda_1^U(p), \cdot)$, where $\lambda_1^U(p)$ denotes firm 1's format strategy conditional on $p' > p$. It can be verified that Proposition 1 extends to the Incumbent-Entrant model. Therefore, $Supp(F_1)$ and $Supp(F_2)$ have the same infimum $p^l < 1$ and the same supremum $p^u = 1$. Therefore, in Nash equilibrium, firm 1's format strategy conditional on $p > p^l$ and firm 2's format strategy conditional on $p < 1$ constitute a Nash equilibrium in the associated hide-and-seek game. These format strategies are equal to the firms' marginal equilibrium format strategies, because as we will verify below, F_1 does not have an atom on p^l and F_2 does not have an atom on $p = 1$.

(ii) Since $p = 1$ is in the support of F_1 and firm 2's format strategy conditional on $p < 1$ max-minimizes v, firm 1's equilibrium payoff is $1 - v^*$. Since firm 1 is chosen with probability one when it charges p^l, it follows that $p^l = 1 - v^*$. But since firm 1's format strategy conditional on $p > p^l$ min-maximizes v, it follows that firm 2's payoff is $v^* \cdot (1 - v^*)$.

(iii) The formulas of F_1 and F_2 follow directly from the condition that every $p \in (1 - v^*, 1)$ maximizes each firm's profit given the opponent's strategy, and the characterization of firm 1's format strategy conditional on $p > p^l$ and firm 2's format strategy conditional on $p < 1$.