

Price Competition under Limited Comparability*

Michele Piccione[†] and Ran Spiegler[‡]

First version: May 2009

This version: December 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 modeling 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. In particular, firms earn max-min payoffs in symmetric equilibria if and only if the graph that represents the comparability between formats satisfies a generalized regularity 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.

1 Introduction

Standard models of market competition assume that consumers rank all the alternatives they are aware of. The ranking may reflect informational constraints, but is complete nonetheless. In reality, consumers are often unable to compare alternatives. Moreover,

*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.

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 discern. 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 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. Price schedules adopted by different firms are often based on different categories, with partly overlapping sections that further complicate the task of comparing them.¹

Marketers and regulators alike 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.² 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...” (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)

¹Of 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.

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

“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 in an abstract, theoretical fashion, 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 interplay between the firms’ pricing and framing decisions? Does greater comparability enhance the consumer’s propensity to switch products?

Our model takes textbook Bertrand competition as a starting point, and adds a notion of comparability as a new dimension. 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 chosen firm, 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 firm 1. Note that, 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 is consistent with the notion that consumers who face complex decision problems are likely to fall back on a default option, if they have one. This casual observation has received experimental support (see, for example, Iyengar and Lepper (2000) and Choi, Laibson and Madrian (2009)), and is consistent with the above-cited consumer protection reports, which emphasize inertia driven by limited comparability as a major cause of low switching rates and weak competitive forces in some 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 .

1.1 An Illustrative Example: “Star” Graphs

We use a simple example to illustrate the model and some of the 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. For simplicity, 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 standard 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 symmetric mixed-strategy Nash equilibrium is unique. When $mq > 1$, the firms’ price and format decisions are 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 a price above p^m . In contrast, when $mq \leq 1$, the firms’ pricing decisions are identical across formats.

The equilibrium structure’s dependence on the core format’s expected number of links is not a coincidence. 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. Therefore, a cheap (expensive) firm has a clear-cut incentive to adopt the core (periphery) as a format strategy. In contrast, when $mq < 1$, each format can induce a higher probability of a price comparison, depending on the rival firm’s format strategy. Notably, the equilibrium format strategy λ^* equalizes the probability of price comparison across formats. As a result, firms are indifferent among all formats, which explains why statistical independence between the firms’ price and format decisions is consistent with equilibrium behavior.

The equilibria in these two regions also differ 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 consistent with consumers’ bounded rationality and

³In 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.

firms' individual rationality. Max-min payoffs can thus be regarded as “constrained competitive profits”. When $mq > 1$, firms earn equilibrium profits above the max-min level. To see why, recall that when firm 1 charges the highest price in the equilibrium distribution, it adopts a peripheral format to minimize 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, which maximizes comparability. In equilibrium, however, whenever firm 2 charges a price above the cutoff 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 are at the max-min level, by a straightforward application of the Minimax Theorem. The fact that the equilibrium format strategy λ^* induces a constant comparison probability implies that λ^* max-minimizes the probability of a price comparison. As a result, when a firm charges the highest price in the equilibrium distribution (which is equal to the consumer's reservation value), it earns max-min profits.

The theoretical implications of the equilibrium analysis for market regulation are somewhat surprising. Current regulatory practice seeks to minimize the number of formats and harmonize them. 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. As to harmonization, suppose that initially, instead of a single major currency there are a number of major currencies, and that the consumer can convert each of them into a minor currency with probability q , and one major currency into another with probability $r \in [q, 1)$. We shall see later that, if a regulator “harmonizes” these major currencies into a single one (as in the original star graph), equilibrium payoffs surprisingly *rise*. Thus, a regulatory intervention that enhances comparability can make the market outcome less competitive, once the firms' equilibrium response to the intervention is taken into account.

1.2 Overview of the Main Results

We begin our analysis of Nash equilibria by providing a simple necessary and sufficient condition for a perfectly competitive market outcome. The rest of our analysis is mostly devoted to symmetric mixed equilibria for order-independent graphs that violate the condition for a competitive outcome. We introduce 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 $mq \leq 1$ cases made in the context of star graphs. A graph is weighted-regular if nodes can be assigned weights such that each node has the same “weighted” number of links to the other nodes. Under weighted regularity, all formats are equally comparable, once the weights are interpreted as the frequency with which they are used. Thus, weighted regularity captures a weak notion of “frame neutrality”.

Our main result is that, in any symmetric equilibrium, firms earn max-min payoffs if and only if the graph that represents the comparability structure satisfies weighted regularity. The significance of this result is that it establishes a tight link between the ability of market forces to push firms to a “constrained competitive” outcome and the potential neutrality of framing. Furthermore, when the graph is weighted-regular, the equilibrium marginal format strategy verifies weighted regularity. In addition, weighted regularity is sufficient and necessary for the property that the equilibrium comparison probability is independent of the firms’ price realizations. In particular, when the graph is not weighted-regular, the firms’ equilibrium price and format decisions are necessarily correlated.

We apply the main result to two special classes of comparability structures. First, we provide a complete characterization for deterministic graphs that represent an equivalence relation. Second, we examine a class of order-independent graphs, referred to as “*bi-symmetric*”, which generalize star graphs. 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 the categories to which they belong. We obtain a closed-form characterization of the (unique) symmetric Nash equilibrium for bi-symmetric graphs. We use this characterization to convey the lesson that regulatory interventions that enhance comparability may have subtle anti-competitive effects. We also show that there is a non-trivial connection between comparability and the extent of consumer switching that occurs in equilibrium.

In later sections, we relax the order-independence assumption and consider an asymmetric variant on our model. The concluding section is devoted to a discussion of the relation between our model and the more conventional view of product differentiation based on preference heterogeneity. Some of the proofs are relegated to an appendix.

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” is somewhat different. 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. And of course, our focus is on market implications of frame dependence rather than on choice-theoretic analysis.

This paper is closely related to Eliaz and Spiegel (2010), which first formalized the idea that framing affects preference incompleteness by influencing the set of alternatives that consumers subject to their preference ranking. There are two major differences. First, Eliaz and Spiegel (2010) mostly interpret a frame as advertising content, and 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 Spiegel (2010), 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 lowering their advertising 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. See Spiegel (2010) for a general treatment of this literature. Within this literature, Spiegel (2006) and Gabaix and Laibson (2006) share the present paper’s preoccupation with firms’ strategic use of “confusing” pricing schemes to increase consumers’ decision errors. In Spiegel (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), Carlin (2009) and Wilson (2010)) stay closer to the rational-consumer paradigm, and model obfuscation as a deliberate attempt to increase consumers’ search costs.

Finally, our paper can be viewed as an extension of a well-known model due to Var-

ian (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. A central feature of our analysis is the characterization of comparability structures (captured by weighted-regular graphs) which give rise to a constant equilibrium comparison probability and yield equilibrium pricing strategies identical to Varian’s.

2 The Model

A *graph* is a pair (X, π) , where X is a finite set of n nodes and $\pi : X \times X \rightarrow [0, 1]$ is a function that determines the probability $\pi(x, y)$ with which a directed edge links node x to node y . Nodes will be referred to 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


We 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. 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 . Thus, firm i ’s payoff under the profile $(p_i, x_i)_{i=1,2}$ is

$$p_i \cdot \frac{1}{2} [1 + \pi(x_j, x_i) \cdot \mathbf{1}(p_i < p_j) - \pi(x_i, x_j) \cdot \mathbf{1}(p_i > p_j)]$$

⁴This assumption is made for expositional simplicity. All our results continue to hold (subject to minor adjustments in the case of Section 5.2) if we assume instead that $\pi(x, x) > 0$ for all $x \in X$.

where $\mathbf{1}$ is the indicator function.

To illustrate the firms' payoff function, consider the graph given by Figure 2, where $X = \{x, y\}$, $\pi(x, y) = q$ 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)

A mixed strategy is a probability measure over $[0, 1] \times X$. We will typically represent mixed strategies as a pair $(\lambda, (F^x)_{x \in \text{Supp}(\lambda)})$, where $\lambda \in \Delta(X)$ is referred to as the (*marginal*) *format strategy*, while F^x is the pricing *cdf* conditional on the format x .

We will also make use of the following pieces of notation. For every interval $I \subset \text{Supp}(F)$, let λ^I denote the format strategy conditional on the event that the price realization lies in an interval I . When I includes only one price p , the conditional format strategy is denoted by λ^p . Given a *cdf* F on $[0, 1]$, let F^- denote its left limit. For any subset non-empty $Z \subseteq X$, $U(Z)$ denotes the uniform distribution over Z .

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

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

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), 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

probability that the seeker finds the hider (or the seeker's payoff) 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 pricing strategies are λ and F , respectively, where the latter is continuous with 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, facing a seeker who plays 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, facing a hider who plays 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

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

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 a format strategy that minimizes the probability of a price comparison.

A firm charging a price that is lower than the opponent's price can enforce a comparison probability of at least v^* , and therefore get a market share of at least $\frac{1}{2}(1 + v^*)$. This is a lower bound on the market share that a firm obtains in any Nash equilibrium when it charges the lowest price in the equilibrium distribution.

2.2 Preliminary Analysis of Nash Equilibria

We will conduct a detailed analysis of Nash equilibria in the sequel. In this sub-section, we present two preliminary findings.

Proposition 1 *The game has a symmetric Nash equilibrium.*

The proof is an application of Corollary 5.3 in Reny (1999), and is omitted. The following proposition gives necessary and sufficient conditions for a competitive equilibrium outcome.

Proposition 2 *In any 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 $y \in X$.*

For the rest of the paper, we assume that the following condition holds:

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

This 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. The proof of Proposition 2 relies on price undercutting arguments that are somewhat subtle. 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, when the original strategy profile is asymmetric, 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 raise the probability of a price comparison when the realization of firm 1's pricing strategy is even lower.

Once a competitive equilibrium outcome has been ruled out, 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.

2.3 Discussion

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

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 $x, y \in X$. 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 behavior. The reason is that the graph represents an intransitive binary relation which induces intransitivity in 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 $[0, 1] \times X$. Both strict and weak preference relations may be intransitive, yet a strict preference relation is acyclic. A probabilistic graph represents a distribution over such incomplete preferences.

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. 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 formats that are hard to compare. 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, since they induce the same payoff function for the firms; they are relevant only for the analysis of switching rates. Only when order independence is relaxed do these assumptions 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 a great deal of sense, given the assumption on default choice in 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 consumer's inability to convert one currency into another could be derived from an earlier decision 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.

As we shall see below, a property of random graphs called "weighted regularity" turns out to be of crucial importance for equilibrium analysis in our model. Therefore, it would be very interesting whether this property is selected or ruled out by a larger model that endogenizes the comparability structure. We leave this question for future work.

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.

At any rate, analyzing the alternative, two-stage 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: they share the payoff function $-\pi$, and, 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.

3 Weighted Regularity

In the next sections, we will focus on the case of order-independent graphs. In this section we introduce a notion of “uniform comparability” across formats, which turns out to be crucial for analysis of Nash equilibria. The degree of comparability of different formats depends on the frequency with which they are adopted, except for the case in which they are all linked with probability equal to one. Whether uniform comparability is potentially obtainable depends on the structure of a graph. Consider, for example, the standard notion of *regularity*: an order-independent graph is regular if there exists a number $\bar{v} > 0$ such that $\sum_{y \in X} \pi(x, y) = \bar{v}$ for all $x \in X$. In a regular graph, the uniform distribution over X induces a uniform comparison probability across all formats. A natural generalization of this notion of potential uniform comparability is obtained by allowing nodes to be chosen with arbitrary probabilities.

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.*

The economic interpretation of weighted regularity is that it is possible for one firm to make its opponent indifferent among all frames - in other words, to “neutralize” the relevance of framing for the rival firm’s competitive strategy. The following are examples of weighted-regular, order-independent graphs.

Example 3.1: Equivalence relations. Consider a deterministic graph 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 $X = \{1, 2, \dots, n\}$, 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. A uniform distribution over all odd-numbered nodes (or all even-numbered nodes) verifies weighted regularity, with $\bar{v} = \frac{2}{n}$.⁵

Example 3.3: Linear similarity. Consider the following deterministic graph. Let $X = \{1, 2, \dots, 3L\}$, 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 $\{3k - 1\}_{k=1, \dots, L}$ verifies weighted regularity.

Example 3.4: Star graphs. 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 $x \neq x_c$:

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

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 $x \neq x_c$.

The following lemma establishes an equivalent definition of weighted regularity, which makes use of the auxiliary hide-and-see game. An order-independent graph is weighted-regular if and only if the associated hide-and-see game has a symmetric Nash equilibrium. To put it somewhat crudely, weighted regularity means that the actions of a firm that maximizes comparability need not be distinct from the actions of a firm that minimizes it.

Lemma 1 *In an order-independent graph (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-see game.*

Proof. (i) Suppose that λ verifies weighted regularity. If one of the players in the associated hide-and-see game plays λ , every strategy for the opponent - including λ itself - is a best-reply. Therefore, λ is a symmetric equilibrium strategy.

(ii) Suppose that (λ, λ) is a Nash equilibrium in the associated hide-and-see game.

⁵Note that any convex combination of these two format strategies also verifies weighted regularity. This is a general property: the set of format strategies that verify weighted regularity for a given graph is convex.

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 . ■

Providing a general characterization of the set of weighted-regular graphs is a difficult problem, which we leave for future work. The following result provides a sufficient condition for weighted regularity.

Proposition 3 *Suppose that for some graph (X, π) , there exists a format strategy $\lambda \in \arg \max \min v$ such that $\lambda(x) > 0$ 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 Nash Equilibrium under Order Independence

In this section, we analyze Nash equilibria for order-independent graphs. We restrict attention to symmetric Nash equilibria. Let $(\lambda, (F^x)_{x \in \text{Supp}(\lambda)})$ be a symmetric Nash equilibrium strategy. Note that the assumption that $\pi(x, x) = 1$ for all $x \in X$ ensures that, by standard arguments, F^x is continuous for any $x \in \text{Supp}(\lambda)$. Therefore, the marginal pricing strategy F is also continuous. In addition, there exists $p^l \in (0, 1)$ such that $\text{Supp}(F) = [p^l, 1]$. This property is entirely conventional in models of imperfect price competition (including Varian (1980)), and the proof is therefore omitted.

We are now ready for the two main results of the paper. First, we establish equivalence between weighted regularity and the property that firms earn max-min payoffs in 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 symmetric equilibrium, each firm's marginal format strategy verifies weighted regularity.*

Proof. (i) *Weighted regularity \implies max-min payoffs.* In fact, we will prove a stronger result. Fix a symmetric Nash equilibrium. For every $p \in [0, 1]$, define $s(p)$ as a firm's

equilibrium market share conditional on charging the price p . We will show that

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

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

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

Thus

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

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

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

By the equilibrium symmetry, each firm's ex-ante market share is equal to $\frac{1}{2}$, and thus (1) 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.

(ii) *Max-min payoffs* \implies *Weighted regularity*. Assume that firms earn max-min payoffs in some symmetric equilibrium. Then, $s(1) = \frac{1}{2}(1 - v^*)$. Recall that $s(p^l) \geq \frac{1}{2}(1 + v^*)$. If this holds with equality, then $\min v(\cdot, \lambda) = \max v(\cdot, \lambda) = v^*$, hence weighted regularity holds. Thus, suppose that $s_i(p^l) > \frac{1}{2}(1 + v^*)$. Since each firm's ex-ante market share is $\frac{1}{2}$, it follows from (2) that there exists a price $p \in (p^l, 1)$ such that

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

Therefore, the following inequality holds for every format strategy $\theta \in \Delta(X)$:

$$\sum_{x \in X} \sum_{y \in X} \theta(x)\lambda(y) \cdot [(1 - 2F^y(p)) \cdot \pi(x, y)] < (1 - F(p))v^* - F(p)v^* \quad (3)$$

This inequality can be rewritten as follows:

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

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 $v^* = \min \max(v)$. Therefore, (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, when firms earn max-min payoffs in equilibrium, 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 implies that the effect of framing on price comparison can potentially be neutralized. The theorem states that the two properties are equivalent: a constrained competitive equilibrium market outcome goes hand in hand with the notion of potential frame neutrality captured by weighted regularity.

For a rough 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 - that is, 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 a graph is not weighted-regular, “acting like a hider” is necessarily distinct from “acting like a seeker”. Therefore, since the marginal format strategy integrates over 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 graph 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 secure a comparison probability of exactly v^* . Thus, the expression $\frac{1}{2}[1 + (1 - F(p))v^* - F(p)v^*]$ effectively serves as a lower bound on the firm’s market share, for any price it considers

charging. 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' 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 that is given by the cdf*

$$F^*(p) = 1 - \frac{1 - v^*}{2v^*} \cdot \frac{1 - p}{p} \quad (4)$$

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

Proof. By Theorem 1, the firms' symmetric equilibrium market share as a function of the price they charge is given by (1). We have established that each firm earns the max-min payoff $\frac{1}{2}(1 - v^*)$ in equilibrium. Therefore, each firm's payoff from any $p \in (p^l, 1)$ is

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

The unique solution to this functional equation is F^* . ■

Equation (4) defines the equilibrium strategy in the two-firm case of Varian's model described in Section 1.3 (Varian (1980)). The intuition for this result is simple. Under weighted regularity the firms' equilibrium market share is determined as a function of the prices they charge, as if the comparison probability was exogenously set at v^* , which is precisely what Varian's model assumes a priori.

Our second main result concerns the equilibrium relation between the probability that consumers make a price comparison and the realizations of the firms' pricing strategies. 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. (i) *Constant comparison probability \implies Weighted regularity.* Assume a constant comparison probability. Then, in particular, $v(\lambda^1, \lambda) = v(\lambda^{p^l}, \lambda)$. But since

$\lambda^1 \in \arg \min v(\cdot, \lambda)$ and $\lambda^{p^l} \in \arg \max v(\cdot, \lambda)$, it follows that λ verifies weighted regularity, such that $v(x, \lambda) = v^*$ for every $x \in X$.

(ii) *Weighted regularity \implies Constant comparison probability.* Recall that in the proof of Theorem 1 (i) we derived equation (1) for every $p \in [p^l, 1]$. This equation can be written as follows:

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

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 the same *as if* it faces a constant comparison probability of v^* . Theorem 2 shows that this is not merely an “as if” property; it does hold in a symmetric equilibrium if and only if the graph 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 a high or a low price - i.e., $\lambda^{p^l} = \lambda^1 = \lambda$. But since λ^{p^l} and λ^1 maximize and minimize $v(\cdot, \lambda)$, respectively, then λ verifies weighted regularity, a contradiction.

However, the converse 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 graph and replicating one of its nodes, so 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 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 3. 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 nodes 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 3)

It should be noted that when a graph is weighted-regular and the hide-and-seek game has a unique 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, $\lambda^{[p^i, p]}$ max-minimizes v and $\lambda^{[p, 1]}$ min-maximizes v . Thus, by hypothesis, $\lambda^{[p^i, p]} = \lambda^{[p, 1]} = \lambda$, and each firm plays the format strategy λ independently of the price it charges.

5 Two Special Cases

In this section we provide a complete analysis of symmetric equilibria for two specific classes of graphs.

⁶This 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 $p = 1$.

5.1 Equivalence Comparability Relations

Example 3.1 presented a class of deterministic graphs that represent an equivalence relation (i.e., a symmetric and transitive binary relation). Such graphs have a special status from a decision-theoretic perspective. Recall that the choice behavior induced by our model is typically inconsistent with utility maximization. Furthermore, the revealed *strict* preference relation is typically intransitive. Among all deterministic graphs, assuming that the graph represents an equivalence relation is equivalent to assuming that the revealed strict preference relation induced by the model of consumer choice is transitive.

As already noted, graphs that represent equivalence relations are weighted-regular. For such graphs, a format strategy verifies weighted regularity if and only if it assigns equal weight to each equivalence class. Let m denote the number of equivalence classes. Then, $v^* = \frac{1}{m}$. By Theorem 2, every symmetric Nash equilibrium exhibits a constant comparison probability of $\frac{1}{m}$. Therefore, it must be the case that in equilibrium, the firms' format strategy assigns probability $\frac{1}{m}$ to each equivalence class conditional on any price in the support.⁷ Thus, we can restate the result as follows. If the revealed strict preference relation induced by the consumer's choice model is transitive, then in Nash equilibrium, firms earn max-min payoffs; furthermore, we can partition the set of formats such that firms mix uniformly over all partition cells independently of the price they charge. This partition can be elicited from individual consumer behavior: x and y belong to the same cell if and only if the consumer displays a strict revealed preference for all pairs (p, x) and (p', y) whenever $p \neq p'$.

5.2 Bi-Symmetric Graphs

In this sub-section, we focus on another special class of graphs, which extends the star graph example 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* $x, y \in X$:

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

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, $n_Z = 1$, $n_Y = m$, and $q_Y = 0$.

⁷When all equivalence classes are singletons, the symmetric equilibrium is then unique. Piccione and Spiegler (2009) prove that there exist no asymmetric equilibria in this case.

Bi-symmetric graphs are attractive because they enable us to capture various “stories” behind limited comparability with simple restrictions on parameter values. 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 are 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 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 $q_Y^* = \frac{1}{m}$, and hence this inequality holds if and only if $m q \leq 1$. When $q_Y^* = q_Z^* = q$, there is a continuum of format strategies that verify weighted regularity. When $(q_Y^* - q)(q_Z^* - q) > 0$, the unique format strategy that verifies weighted regularity assigns probability

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

to the set Z , and mixes uniformly within each of the sets Y and Z . We denote this format strategy by λ^* . In this case, the hide-and-seek game has (λ^*, λ^*) as the unique Nash equilibrium.

The value of the hide-and-seek game under weighted regularity is

$$v^* = \frac{q_Y^* q_Z^* - q^2}{(q_Y^* - q) + (q_Z^* - q)} \quad (5)$$

when $(q_Y^* - q)(q_Z^* - q) > 0$, and $v^* = q$ when $(q_Y^* - q)(q_Z^* - q) = 0$.

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

Proposition 4 *Let (X, π) be a bi-symmetric graph. If $(q_Y^* - q)(q_Z^* - q) > 0$, there is a unique symmetric Nash equilibrium, in which firms play the format strategy λ^* , and the*

pricing strategy (4) at every node, where v^* is given by (5). Firms earn the max-min payoff $\frac{1}{2}(1 - v^*)$.⁸

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 $\mathcal{U}(Z)$ ($\mathcal{U}(Y)$). We use this observation to construct a symmetric equilibrium strategy which has 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 $\mathcal{U}(Z)$, and the format strategy conditional on any price $p \in (p^m, 1]$ is $\mathcal{U}(Y)$. The marginal pricing strategy F satisfies:

$$F(p^m) = \frac{q - q_Y^*}{q_Z^* - q_Y^*} \quad (6)$$

Note that the total probability that the marginal format strategy assigns to the set Z (Y) is $F(p^m)$ ($1 - F(p^m)$). The conditional pricing strategies are given by the following pair of functional equations which constitute the indifference conditions that characterize the mixed-strategy equilibrium. Let F^Z (F^Y) denote the pricing *cdf* conditional on playing a format in Z (Y). For every $p \in [p^l, p^m]$:

$$\frac{p}{2}[1 + F(p^m)(1 - 2F^Z(p))q_Z^* + (1 - F(p^m))q] = \frac{1}{2}[1 - F(p^m)q - (1 - F(p^m))q_Y^*] \quad (7)$$

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

$$\frac{p}{2}[1 + (1 - F(p^m))(1 - 2F^Y(p))q_Y^* - F(p^m)q] = \frac{1}{2}[1 - F(p^m)q - (1 - F(p^m))q_Y^*] \quad (8)$$

The R.H.S on each of these two equations represents the firms’ equilibrium payoff. Plugging in (6), we obtain:

$$\frac{1}{2}\left[\frac{q - q_Y^*}{q_Z^* - q_Y^*}(1 - q) + \frac{q_Z^* - q}{q_Z^* - q_Y^*}(1 - q_Y^*)\right] \quad (9)$$

Observe that this expression for the firms’ equilibrium payoff 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.⁹

⁸When $q_Y^* = q_Z^* = q$, the result is slightly weaker. In symmetric equilibrium, the marginal framing strategy verifies weighted regularity, and the pricing strategy is (4), where $v^* = q$. However, the infinite number of framing strategies that verify weighted regularity can give rise to payoff-irrelevant correlation between the firms’ pricing and framing decisions.

⁹To check that the strategy given by (6)-(8) is indeed a symmetric equilibrium strategy, all we need

Proposition 5 *Let (X, π) be a bi-symmetric graph. If $(q_Y^* - q)(q_Z^* - q) < 0$, there is a unique symmetric Nash equilibrium, which is the cutoff equilibrium characterized by (6)-(8). The firms' equilibrium payoff is given by (9).*

As mentioned before, the classification of bi-symmetric graphs into those that satisfy weighted regularity and those that do not matches two different interpretations of the set of formats Y 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, 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.

6 Does Greater Comparability Lead to a More Competitive Outcome?

A basic intuition that underlies the consumer protection statements quoted in the Introduction is that greater comparability of price formats makes the market more competitive and therefore favors consumers. Indeed, if consumers faced a fixed set of price-format pairs, switching from a comparability structure π to another structure π' that satisfies $\pi'(x, y) \geq \pi(x, y)$ for every $x, y \in X$ would make consumers weakly better off, because the probability they will choose the cheapest alternative can only go up.

Is the competitive effect of greater comparability robust to equilibrium analysis? When π' is weighted-regular, the answer is clearly 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 π' is not weighted-regular. Consider the case of bi-symmetric graphs that violate weighted regularity, where equilibrium payoffs are given by (9). 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

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.

and q_Y^*), 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_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 Y formats. 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 enjoy greater market power *de facto*. We can see that “local” improvements in comparability - such as increasing the transparency of already-simple formats - may have a detrimental impact on consumer welfare.

7 Consumer Switching

The case of bi-symmetric graphs also enables us to address the issue of consumer switching, and qualify the message of the consumer protection reports quoted in the Introduction, namely that greater comparability leads to more frequent switching.

In 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 $\frac{1}{2}$. This is an immediate corollary of the symmetry of π : conditional on making a comparison, the consumer faces a symmetric posterior probability distribution over price profiles (p_1, p_2) , independently of the identity of the consumer’s default option. 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. Under weighted regularity, we saw that the comparison probability is v^* , independently of the prices that firms charge, and therefore the switching rate is $\frac{1}{2}v^*$ in equilibrium. Thus, when we compare two weighted-regular graphs, 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. In the case of

bi-symmetric graphs, the equilibrium comparison probability 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 the firms' equilibrium price and format decisions are correlated, the positive link between the switching rate and market competitiveness may break down.

8 Order-Dependent Graphs

In this section we relax order independence and explore the robustness of our main result to this extension. We begin by extending the notion of weighted regularity.

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

Observe that this definition is reduced to weighted regularity when the graph is order-independent. The following example illustrates the difference between the two concepts. Let $X = \{a, b, c\}$, $\pi(a, b) = \pi(a, c) = 1$ and $\pi(x, y) = 0$ for all other distinct x, y . A format strategy that assigns probability $\frac{1}{2}$ to each of the formats b and c verifies weak weighted regularity. However, the graph is not weighted-regular.

When β verifies weak weighted regularity, the two weak inequalities in Definition 2 are binding for every x in the support of β . To see why, suppose that the L.H.S inequality in Definition 2 is strict for some $x \in X$ for which $\beta(x) > 0$. Then, summing over all $x \in X$, we obtain

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

Now consider the R.H.S inequality in Definition 2. If we sum over all $x \in X$, we obtain

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

a contradiction.

Building on this observation, it is possible to establish an equivalent definition of weak weighted regularity in terms of the associated hide-and-seek game, as in Section 3. This equivalence implies that when β verifies weak weighted regularity, \bar{v} is equal to v^* , the value of the associated hide-and-seek game. The proof is omitted because it proceeds as the proof of the analogous Lemma 1.

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

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

Proposition 6 *Suppose that (X, π) satisfies weak weighted regularity. Then, firms earn max-min payoffs in any symmetric Nash equilibrium.*

The proof follows the same line of reasoning as Theorem 1(i). Fix a symmetric Nash equilibrium. By weak weighted regularity, there exists 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^* ; 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

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

This is exactly the same lower bound we obtained under order-independence, and the proof proceeds in the same manner.

The converse to this result does not hold in general. When an order-independent graph violates weak weighted regularity, it does not follow that firms necessarily earn payoffs above the max-min level in symmetric equilibrium. For example, recall the graph given by Figure 2: $X = \{x, y\}$, $\pi(x, y) = q$ and $\pi(y, x) = 0$. this graph violates weak weighted regularity. However, it admits a symmetric Nash equilibrium in which 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 $[\frac{1-q}{3-q^2}, \frac{1}{3+q}]$. 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 in this equilibrium.

9 Asymmetric Firm Assignment

Equilibrium analysis under order dependence is greatly simplified if we drop the assumption that the consumer’s initial firm assignment is symmetric. 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 $(\lambda_i, (F_i^x)_{x \in \text{Supp}(\lambda_i)})_{i=1,2}$ 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:

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

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 Incumbent’s (Entrant’s) unequivocal incentive to avoid (foster) price comparisons. Each firm acts as if it has a fixed role in the hide-and-seek game, independently of the price it charges. This implies that their pricing decisions are made as if they play an asymmetric version of the Varian model.

10 Concluding Remarks

This paper studied 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 (2010), 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 .

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. In conventional models product differentiation is viewed as the market's response to consumers' differentiated tastes.

To understand the comparison 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. In our model consumer choice is typically impossible to rationalize with a random utility function over pairs (p, x) , whereas 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, 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 their conventional spatial-competition analogue.

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. Conversely, 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 other hand, 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 $u_\theta : X \rightarrow \{0, 1\}$. The function u_θ essentially describes the set of brands that type θ likes, whereas the graph π_θ determines the type's ability to compare different brands. Exploring this model, and particularly its ability to account for patterns of consumer behavior, 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. (2009): "Strategic Price Complexity in Retail Financial Markets," *Journal of Financial Economics* **91**, 278-287.
- [3] Chioveanu, I. and J. Zhou (2009): "Price Competition with Strategic Obfuscation", UCL-ELSE Working Paper no. 339/2009.

- [4] 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.
- [5] 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.
- [6] Ellison, G. and A. Wolitzky (2008): “A Search Cost Model of Obfuscation”, mimeo. MIT.
- [7] Eliaz, K. and R. Spiegler (2010): “Consideration Sets and Competitive Marketing,” *Review of Economic Studies*, forthcoming.
- [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] OFT (2008): “Personal Current Accounts in the UK”, an OFT market study, http://www.oft.gov.uk/shared_oftr/reports/financial_products/OFT1005.pdf.
- [11] Piccione, M. and R. Spiegler (2009): “Framing Competition,” UCL-ELSE Working Paper no. 336/2009.
- [12] Reny P. (1999), On the Existence of Pure and Mixed Strategy Equilibria in Discontinuous Games, *Econometrica* **67**, 1029-1056.
- [13] Rubinstein, A. and Y. Salant (2008): “(A,f): Choices with Frames,” *Review of Economic Studies*.**75**, 1287-1296.
- [14] Spiegler R. (2006): “Competition over Agents with Boundedly Rational Expectations,” *Theoretical Economics* **1**, 207-231.
- [15] Spiegler R. (2010): *Bounded Rationality and Industrial Organization*. Forthcoming in Oxford University Press, New York.
- [16] Varian, H. (1980): “A Model of Sales,” *American Economic Review* **70**, 651-59.

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

11 Appendix: Proofs

11.1 Proposition 2

Define $X^A = \{x \in X : \pi(y, x) = 1 \text{ for all } y \in X\}$. Suppose that $F_1(0) = F_2(0) = 1$. Then, both firms earn zero profits. If $\lambda_i(x) > 0$ and $\pi(y, x) < 1$ for some $x \in \text{Supp}(\lambda_i)$ and some $y \in X$, then firm j can make positive profits charging $p = 1$ and choosing y , a contradiction. It follows that $\text{Supp}(\lambda_i) \subseteq X^A$, hence X^A is non-empty.

Suppose now that X^A is non-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 $x \in \text{Supp}(\lambda_2)$, $y \in \text{Supp}(\lambda_1)$. Assume the contrary. Let \bar{p}_i denote the supremum of $\text{Supp}(F_i)$, and denote $\bar{p} = \max(\bar{p}_1, \bar{p}_2)$. Without loss of generality, assume $\bar{p} = \bar{p}_2$. Take a node z in the support of λ_2 such that $\bar{p} \in \text{Supp}(F_2^z)$. Firm 2's profit is

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

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

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

Since firm 2's payoff is positive, $F_1^{x^-}(\bar{p}) < 1$ for some $x \in \text{Supp}(F_1)$. 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 $\text{Supp}(F_1) \cup \text{Supp}(F_2)$ for which there exist $x \in \text{Supp}(\lambda_j)$ and $y \in \text{Supp}(\lambda_i)$, $i \neq j$, such that $p \in \text{Supp}(F_i^y)$ and $\pi(x, y) < 1$. If $p^* = p^l$, then for any $y' \in X^A$, the pure strategy $(p^l - \varepsilon, y')$ outperforms the pure strategy (p^l, y) , for ε sufficiently small, a contradiction. Therefore, $p^* > p^l$. Without loss of generality, suppose that $p^* \in \text{Supp}(F_2^y)$. Firm 2's payoff from the pure strategy (p^*, y) is

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

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

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

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

11.2 Proposition 3

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 $\mu(x) > 0$ for all $x \in X$, there exists $\tilde{\mu} \in \Delta(X)$ such that, for all $y \in X$,

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

Order the nodes so that $X = \{1, \dots, n\}$. Any $\beta \in \Delta(X)$ is thus represented by a row vector $(\beta_1, \dots, \beta_n)$. Let Π be a $n \times n$ matrix whose ij th entry is $\pi(i, j)$. 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 $\Pi\beta^T = (c, c, \dots, c)^T$. By Farkas' Lemma, there exists a column vector $\delta \in \mathbb{R}^n$ such that $\Pi\delta \geq 0$ and $(c, c, \dots, c)\delta < 0$. Since $\pi(i, i) = 1$ for every $i \in \{1, \dots, n\}$ and $\pi(i, j) \geq 0$ for all $i, j \in \{1, \dots, n\}$, 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 $i \in \{1, \dots, n\}$. 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, $\lambda(x) > 0$ 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 . This contradicts the assumption that λ is a max-min strategy for the seeker.

11.3 Proposition 5

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

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

Observe that by the assumption that the graph violates weighted regularity, $ab < 0$.

Let $(\lambda, (F^x)_{x \in \text{Supp}(\lambda)})$ 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^{x^l} and p^{x^u} 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,

$$v^x(\lambda) = \sum_{y \in X} \lambda(y) \pi(x, y) \quad (10)$$

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 proof relies on a series of lemmas.

Lemma 3 $\lambda(x) = \lambda(x')$ 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

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

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

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

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

Lemma 4 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

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

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

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

By Lemma 3, $\lambda(y) = \lambda(y')$ and therefore $v^y(\lambda) = v^{y'}(\lambda)$. It follows that the deviation is profitable. ■

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

Proof. Suppose that $\lambda(x) = 0$ for some $x \in Y$. By Lemma 3, 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, a contradiction. ■

Lemma 6 For any $y \in Y$ and $z \in Z$, $p^{yu} = p^{zl}$ or $p^{zu} = p^{yl}$.

Proof. Suppose that $v^z(\lambda) < v^y(\lambda)$. By Lemma 4, the nodes in Y have the same F^y and the nodes in Z have the same F^z . Therefore, $S^y \cap S^z \neq \emptyset$, for any $y \in Y$ and $z \in Z$. The following equations must hold in equilibrium.

$$\begin{aligned} \lambda(z) q n_Z (1 - F^z(p^{yu})) + \frac{1}{2} (1 - v^y(\lambda)) &= \\ \lambda(z) (1 + q_Z (n_Z - 1)) (1 - F^z(p^{yu})) + \frac{1}{2} (1 - v^z(\lambda)) & \\ \lambda(z) q n_Z + (1 + q_Y (n_Y - 1)) \lambda(y) ((1 - F^y(p^{zl}))) + \frac{1}{2} (1 - v^y(\lambda)) &= \\ \lambda(z) (1 + q_Z (n_Z - 1)) + q n_Y \lambda(y) ((1 - F^y(p^{zl}))) + \frac{1}{2} (1 - v^z(\lambda)) & \end{aligned}$$

which simplify to

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

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

By Lemmas 5 and 6, a symmetric Nash equilibrium must be a cutoff equilibrium. Moreover, by Lemma 4, it suffices to consider two cases: either $\lambda^{[p^m, 1]}$ is a uniform distribution over Y and $\lambda^{[p^l, p^m]}$ is a uniform distribution over Z , or $\lambda^{[p^m, 1]}$ is a uniform distribution over Z and $\lambda^{[p^l, p^m]}$ 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

$$\lambda(y)n_Yq - \lambda(z)n_Zq_Z^* = \lambda(y)n_Yq_Y^* - \lambda(z)n_Zq$$

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

$$\lambda(z)n_Zq - \lambda(y)n_Yq_Y^* = \lambda(z)n_Zq_Z^* - \lambda(y)n_Yq$$

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 .

11.4 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 formats minimize $v(\cdot, \lambda_2^{[p^l, p]})$. Similarly, for every price p that lies strictly below the supremum of $Supp(F_1)$, firm 2's optimal formats maximize $v(\lambda_1^{[p, 1]}, \cdot)$. By standard arguments, the closure of both $Supp(F_1)$ and $Supp(F_2)$ is $[p^l, 1]$, where $p^l > 0$. Therefore, 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 $p < 1$.