



HAL
open science

Nonlinear Pricing in Oligopoly: How Brand Preferences Shape Market Outcomes

Renato Gomes, Jean-Marie Lozachmeur, Lucas Maestri

► **To cite this version:**

Renato Gomes, Jean-Marie Lozachmeur, Lucas Maestri. Nonlinear Pricing in Oligopoly: How Brand Preferences Shape Market Outcomes. 2022. hal-03629496

HAL Id: hal-03629496

<https://hal.science/hal-03629496>

Preprint submitted on 4 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

March 2022

“Nonlinear Pricing in Oligopoly: How Brand Preferences
Shape Market Outcomes”

Renato Gomes, Jean-Marie Lozachmeur and Lucas Maestri

Nonlinear Pricing in Oligopoly: How Brand Preferences Shape Market Outcomes*

Renato Gomes[†] Jean-Marie Lozachmeur[‡] Lucas Maestri[§]

First Version: 23/04/2019

This Version: 21/01/2022

Abstract

We study oligopolistic competition by firms practicing second-degree price discrimination. In line with the literature on demand estimation, our theory allows for comovements between consumers' taste for quality and propensity to switch brands. If low-type consumers are sufficiently less (more) brand loyal than high types, (i) quality provision is inefficiently low at the bottom (high at the top) of the product line, and (ii) informational rents are negative (positive) for high types, while positive (negative) for low types. We produce testable comparative statics on pricing and quality provision, and show that more competition (in that consumers become less brand-loyal) is welfare-decreasing whenever it tightens incentive constraints (so much so that monopoly may be welfare-superior to oligopoly). Interestingly, pure-strategy equilibria fail to exist whenever brand loyalty is sufficiently different across consumers types. Accordingly, price/quality dispersion ensues from the interplay between self-selection constraints and heterogeneity in brand loyalty.

JEL Classification: D82

Keywords: competition, price discrimination, asymmetric information, preference correlation, price dispersion

*A previous version of this paper was circulated under the title “Brand Substitutability and Quality Provision in Oligopolistic Price Discrimination.” For very helpful feedback, we thank, without implicating, Bruno Jullien, Patrick Rey, Nicolas Schutz, Jeroen Swinkels and Mariya Teteryatnikova, as well as seminar participants at the Universities of Bologna, Glasgow, Edinburgh, Essex, Mannheim, Zurich, Bocconi, Collegio Carlo Alberto, FGV-EPGE, the Higher School of Economics, the European Commission, the Lisbon Meetings of Game Theory and Applications, EARIE 2020, the CEPR Virtual IO Seminar, and the CSIO-TSE 18th Joint Workshop on Industrial Organization. We are grateful to Céline Parzani for superb research assistance. Renato Gomes gratefully acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 759733 - PLATFORM). Renato Gomes and Jean-Marie Lozachmeur also acknowledge funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir programme).

[†]Toulouse School of Economics, CNRS, University of Toulouse Capitole: renato.gomes@tse-fr.eu.

[‡]Toulouse School of Economics, CNRS, University of Toulouse Capitole: jean-marie.lozachmeur@tse-fr.eu.

[§]FGV-EPGE, Escola Brasileira de Economia e Finanças: maestri.lucas@gmail.com.

1 Introduction

Since as early as 1849, with the pioneering work of Jules Dupuit, economists have investigated the effects of second-degree price discrimination on pricing and quality provision. Anticipating the modern treatments of Mussa and Rosen (1978) and Maskin and Riley (1984), Dupuit noted that profit maximization by a monopolist leads to under-provision of quality at the bottom of the product line. Intuitively, under-provision prevents consumers with high valuations per quality from purchasing low-quality products, therefore making it possible to set high prices at the top of the product line, where more profits can be made.

This clarity of insight is missing in oligopolistic settings, where at least two firms compete for consumers by offering menus of products. One critical feature of these markets is that consumers's preferences over product characteristics are often correlated with their propensity to switch brands. The latter possibility has been recognized for long, and plays a key role in the empirical literature that estimates demand employing discrete-choice models with random coefficients (in the tradition of Berry, Levinsohn and Pakes 1995). In such models, consumers choose among different options from competing product lines by weighing their respective price and quality dimensions. To produce more flexible estimates, it is often assumed that consumer preferences over price/quality attributes are random, while depending on demographics such as income, age, family size, etc (see, for instance, Nevo 2001). These studies often find that consumers' price sensitivity and taste for quality are correlated, which implies that consumer segments (along the product line) systematically differ on their propensities to switch brands (e.g., in response to price discounts).¹

While there is no reason to expect this correlation to be always positive or negative (across markets), it is intuitive that common factors determine both the consumers' tastes for quality and brand loyalty. To illustrate, suppose income is the main factor behind one's tastes for quality (e.g., high earners like premium products) as well as behind brand switching (e.g., those with higher marginal utility of money are more likely to react to price discounts). In this case, consumers with stronger tastes for quality (high earners) are less likely to switch brands (for having a lower marginal utility of money). This is consistent with the empirical findings of Kaplan and Menzio (2015) and Kaplan et al (2019), who show that consumers with higher incomes are more likely to remain brand-loyal, while caring more about product quality.² This is also consistent with Petrin (2002), who finds, in the context of minivans, that those consumers with weaker brand preferences assign less value to quality attributes such as horsepower or vehicle size.

¹This is particularly relevant for structural empirical work investigating product design in oligopolistic settings. Examples include Gandhi et al. (2008), Chu (2010) and Fan (2013).

²Anecdotal evidence suggests that a related story applies to airline markets. For instance, the most loyal customers of Air France are likely to be business travelers who simultaneously exhibit higher tastes for premium features (extra leg space, access to the pre-boarding lounge, etc), as well as a lower propensity to switch airlines (be it because the employer pays the ticket, so they are less price sensitive, or because of the convexity of frequent flier rewards).

By contrast, Crawford, Scherbakov and Shum (2019), in the context of cable TV, find that consumers with stronger tastes for quality are easier to poach through price discounts than those who exhibit weaker tastes for quality.³ One possible explanation to this finding is that consumers who purchase premium packages (defined as containing more channels) are those who spend more time watching TV, earn lower incomes,⁴ and therefore exhibit a higher marginal utility of money (being therefore more price-sensitive).⁵ The authors then show that the quality of premium packages is set inefficiently high by cable companies, which contrasts with the received wisdom from the monopolistic screening literature (and also from oligopolistic models, as reviewed below).

All in all, the applied literature, as well as casual observations, suggest that correlation in consumer preferences is not only empirically relevant, but also consequential for pricing and product design. Yet, theory is mostly silent about how comovements between consumers' tastes for quality and brand loyalty affect market outcomes under competition. We try to fill this gap.

Model and Results

We embed the canonical Mussa and Rosen (1978) model of price discrimination into a discrete-choice framework, with an arbitrary number of firms. Crucially, the dispersion of brand tastes (which determines the amount of brand switching in response to price discounts) is assumed to vary with the consumer's taste for quality. Accordingly, "high types" (i.e., those consumers with a high valuation for quality) may be more or less loyal to their preferred brand than "low types." As we shall see, this degree of flexibility is crucial to explain the diversity of market outcomes observed under competition.

Firms simultaneously offer menus of price-quality pairs designed to screen consumers' unobserved tastes for quality. Firms are blind to consumer brand preferences, reflecting the anonymity of past market transactions, or privacy regulation.

To fix ideas, we first revisit the case of a "balanced oligopoly," where the consumers' propensity to switch brands is independent of their tastes for quality. In line with the seminal works of Armstrong and Vickers (2001) and Roche and Stole (2002), we find that quality provision is efficient in equilibrium provided consumers are not too brand loyal. We however depart from these contributions, rather following Bénabou and Tirole (2016), on how to model consumers' participation decisions. Crucially, in our model, variations in consumers' propensity to switch brands do not affect participation, enabling us to cover the whole spectrum of competitive intensity. Accordingly, we

³A similar pattern is found by Durrmeyer (2020), who studies the automobile market in France. She finds that consumers with stronger preferences for "green" attributes (favoring cars which fuels emit less CO₂) are the most sensitive to automobile prices.

⁴There is indeed robust empirical evidence showing that the time spent watching TV is negatively correlated with household income. See for instance Nielsen (2015).

⁵Another possibility is suggested by the behavioral literature on rational inattention (see Gabaix 2014 and the references therein). This literature argues that consumers who spend more money on a product (premium cable TV) tend to be more attentive to its price, and therefore more likely to switch brands in response to price differences.

show that, as consumers become more loyal to their preferred brands, the equilibrium approaches the monopolistic outcome of Mussa and Rosen (1978).⁶

In the general case where brand loyalty is type-specific, our analysis delivers four main insights. First, relative to its efficient level, equilibrium menus over-provide quality at the top of the product line if the propensity of low-type consumers to switch brands is small relative to that of high types. Intuitively, under this form of preference correlation, firms enjoy more market power among low types, who then obtain lower payoffs in equilibrium. To avoid profit dissipation (stemming from low types selecting the premium product, which profit margin is smaller), firms then inefficiently raise the quality of the premium product. This is consistent with the aforementioned contribution of Crawford, Scherbakov and Shum (2019), who estimate “low types” to be less prone to switch cable companies, while finding that cable companies design premium packages of inefficiently high quality.⁷

Conversely, equilibrium menus under-provide quality at the bottom of the product line if the propensity of high-type consumers to switch brands is small relative to that of low types. The intuition is the mirror image of that from the previous case: Here, firms enjoy more market power among high types, who then obtain lower payoffs in equilibrium. To avoid profit dissipation (stemming now from high types selecting the baseline product), firms then inefficiently reduce the quality of the baseline product. This prediction is consistent with McManus (2007), who finds that (oligopolistic) coffee shops distort product sizes for “sweet espresso” drinks, choosing inefficiently small servings except at the largest cup size.⁸ It is also consistent with anecdotal evidence from airline services: lack of comfort is common in economy class seats, which are typically foregone by business travelers exhibiting more brand loyalty than cheapskate tourists.

Second, we show that asymmetric information about one’s tastes for quality may either benefit or hurt consumers, depending on the correlation between preferences for quality and brand loyalty. To understand the novelty of this finding, let us reconsider the monopolist benchmark of Mussa and Rosen (1978). Because of the self-selection constraints inherent to price discrimination, all types (weakly) benefit from privately knowing their tastes for quality (i.e., informational rents are necessarily non-negative). This conclusion holds true under competition if the consumers’ brand loyalty is independent of their tastes for quality. The reason is that low types obtain zero payoffs whenever incentive constraints bind, which implies they are indifferent between the cases of complete and asymmetric information about their preferences. In turn, high types gain even more from asymmetric information as one moves from monopoly to oligopoly. The reason is that asymmet-

⁶In contrast to our paper, Bénabou and Tirole (2016) study competition in linear contracts in a common-value environment.

⁷See also Crawford (2012) for an earlier discussion on how to measure quality distortions in empirical models of differentiated product demand.

⁸McManus (2007) posits in his structural model that price sensitivity is constant, being therefore orthogonal to (random) preferences over product attributes. It is likely, though, that consumers of “sweet espresso drinks,” which are the most differentiated across shops, have strong brand preferences, the more so among those who drink more coffee.

ric information magnifies competition, as relinquishing more utility to high types relaxes incentive constraints, increasing the efficiency of low-type contracts.

By contrast, informational rents are negative to high types (but positive to low types) if brand loyalty is higher among low types. The reason is that, under this form of preference correlation, asymmetric information weakens competition for high types, as relinquishing more utility to these consumers tightens incentive constraints, decreasing the efficiency of premium products. Conversely, increasing the indirect utility of low types alleviates the upward distortion at the top of the product line, which intensifies competition for these consumers. As a result, private information about one's tastes benefits low types but hurts high types.

On the other hand, informational rents are positive to high types (but negative to low types) if brand loyalty is higher among high types. The reason is that, under this form of preference correlation, asymmetric information weakens competition for low types, as relinquishing more utility to these consumers tightens incentive constraints, decreasing the efficiency of baseline products. Conversely, increasing the indirect utility of high types alleviates the downward distortion at the bottom, which intensifies competition for these consumers. Therefore, private information about one's tastes benefits high types but hurts low types. Under either form of preference correlation, these results contrast with the received wisdom according to which consumers are necessarily better off under second- rather than third-degree price discrimination (where pricing by firms is not constrained by consumer self-selection) - see, for instance, Varian (2006).

Third, we develop a number of comparative statics on pricing and quality provision that eluded previous analysis. For instance, we find that, as low (resp., high) types become more prone to switch brands, quality provision increases (resp., decreases) along the product line. As a result, welfare decreases as low types become more prone to switch brands if quality provision is excessive at the top, while it increases as low types become more prone to switch brands if quality provision is deficient at the bottom of the product line. In the former case (over-provision at the top), this effect is so pronounced that monopoly may produce higher welfare (although lower consumer surplus) than oligopoly. Relatedly, we also show that the price charged to low types is non-monotone in the brand loyalty of high-type consumers. These implications are testable, and differentiate our model from other theories of price discrimination under competition.

Fourth, we show that pure-strategy equilibria fail to exist whenever brand loyalty is sufficiently different across consumers types, which implies equilibria are necessarily in mixed strategies. This non-existence result is driven by the interplay between self-selection constraints and the fact that different types exhibit different propensities to switch brands. Accordingly, our theory identifies a new rationale for price/quality dispersion in private-value settings, unlike previous literature that relates dispersion to search/informational frictions, or to the unobservability of inventories.⁹ Importantly,

⁹See Varian (1980), Burdett and Judd (1983) and Baye and Morgan (2001) for classical treatments of pricing under

we also characterize mixed-strategy equilibria, producing new (testable) predictions about the distribution of market offers. We cautiously interpret this result as consistent with the fact that many oligopolistic markets practicing second-degree price discrimination are “unstable,” in that product features and prices are constantly revised by competing firms (e.g., air travel).¹⁰

Paper Outline

The rest of the paper is organized as follows. Section 2 sets up the model. Section 3 derives some preliminary results, and revisits the benchmark where consumers’ propensity to switch brands is independent of tastes for quality. Section 4 considers the general case where brand loyalty is type-dependent, characterizing pure-strategy equilibria and deriving comparative statics. Section 5 explains why firms’ offers might be dispersed in equilibrium, and studies mixed-strategy equilibria. Section 6 discusses some modeling assumptions and extensions to the baseline model, including the case of a continuum of types. Section 7 reviews the pertinent literature, while section 8 collects our main empirical implications. Proofs are in the Appendix at the end of the document.

2 Model

There is a unit-mass continuum of consumers with single-unit demands, and $J \geq 2$ firms indexed by $j \in \{1, \dots, J\}$. Consumers are heterogeneous in their tastes for quality, denoted by θ , and their tastes for brands, described by the vector $\varepsilon \equiv (\varepsilon^1, \dots, \varepsilon^J)$. For each consumer, θ is a draw from a distribution with binary support $\{\theta_l, \theta_h\} \subset \mathbb{R}_{++}$, where $\Delta\theta \equiv \theta_h - \theta_l > 0$, and associated probabilities p_l and p_h (with $p_l, p_h > 0$ and $p_l = 1 - p_h$).¹¹ In turn, the random vector ε is a draw (independent of θ) from a symmetric distribution G with convex support contained on \mathbb{R}^J and density g .¹² The pair (θ, ε) is private information of each consumer, and independently drawn across consumers. For convenience, we abuse terminology and refer to the quality taste θ as the consumer’s *type*.¹³

Each firm j designs a menu (of arbitrary size) of quality-price pairs (or contracts), indexed by $n \in N^j$, where N^j is the set of indexes associated with firm j ’s menu. We denote by q_n^j the quality and by y_n^j the price of contract n offered by firm j .¹⁴ To guarantee full market coverage, we model consumers’ (non-)participation in the following way: we add a “no-purchase contract” $(0, 0)$ to the

search/informational frictions, and Shelegia and Wilson (2020) for a substantial generalization of previous results. On how unobservable inventories and advance production generate price dispersion, see Montez and Schutz (2020).

¹⁰Some of these markets do not suffer from severe informational frictions, nor involve advance production, but do seem to exhibit heterogeneity in consumers’ brand loyalty.

¹¹See subsection 7.1 for the case of a continuum of types.

¹²We further impose some weak regularity conditions; namely, that the density g is differentiable in the interior of the support and that ε^j has finite moments for all j .

¹³This terminology reflects the fact that firms cannot screen consumers’ tastes for brands.

¹⁴We thus rule out stochastic as well as reciprocal mechanisms (where the offer of a firm may depend on that of its competitor).

menu of each firm, while assuming that consumers must pick a contract in some menu. Intuitively, consumers can always find a low-quality low-price outside good that shares each firm j 's taste shock (see Remark 1 and subsection 6.3 for further discussion). Accordingly, we consider menus of the form $m^j \equiv \{(q_n^j, y_n^j) : n \in N^j\}$, with the understanding that $0 \in N^j$ and that $(q_0^j, y_0^j) = (0, 0)$ is firm j 's no-purchase contract.

The utility that a consumer with taste vector ε and taste for quality θ_k , $k \in \{l, h\}$, obtains from purchasing from firm j is then

$$\max_{n \in N^j} \{\theta_k q_n^j - y_n^j\} + t_k \varepsilon^j, \quad (1)$$

where the *brand loyalty parameter* $t_k \geq 0$ captures the intensity of brand preferences by type- k consumers. The comparison between t_l and t_h determines whether consumers with low or high tastes for quality are more prone to switch brands in response to changes in firms' offers. Naturally, the brand loyalty parameter depends on the type of the consumer, but not on the contract chosen.

Firms incur a per-unit cost $\varphi(q)$ for providing a good of quality q . The cost function $\varphi(\cdot)$ is twice continuously differentiable, strictly increasing, strictly convex, and satisfies $\varphi(0) = \varphi'(0) = 0$ and $\lim_{q \rightarrow \infty} \varphi'(q) = \infty$. Accordingly, the profit *per sale* by firm j of its contract n is $y_n^j - \varphi(q_n^j)$. Firm j 's profit from contract n then equals the demand for this contract times its profit per sale. The firm's total profit adds up profits across all contracts in its menu.

Firms simultaneously post menus, after which each consumer chooses her preferred contract across firms' menus. A (possibly mixed) strategy by each firm is a distribution over menus σ^j . A symmetric equilibrium (for short, equilibrium), possibly in mixed strategies, is a distribution over menus σ^* that is a best response to itself.

3 Preliminaries

3.1 A Change of Variables

Because types are binary in the baseline model, it is without loss of generality to assume that firms offer menus with at most two contracts (beyond no purchase). It is then useful to identify each contract n offered by firm j with the consumer type θ_k that would select that contract, *conditional* on choosing firm j . Our choice of labels then implies the following incentive-compatibility constraints for each $j \in \{a, b\}$ and $k \in \{l, h\}$:

$$IC_k^j : \quad u_k^j \equiv \theta_k q_k^j - y_k^j \geq \theta_k q_n^j - y_n^j \quad \text{for all } n \in \{0, l, h\}.$$

We refer to u_k^j as type- k 's indirect utility under firm j 's menu (gross of brand preferences). In particular, condition IC_k^j implies that the contract targeting each type k has to be weakly superior to the no-purchase contract. Therefore, $u_k^j \geq 0$, which is the customary individual rationality constraint.

In light of the above, we can compare the net utilities in (1) to derive the firms' demands from each consumer type. To do so, let $\mathbf{u}_k^{-j} \equiv (u_k^1, \dots, u_k^{j-1}, u_k^{j+1}, \dots, u_k^J)$ and denote by

$$H_k^{j,l} \left(x \mid \mathbf{u}_k^{-j} \right) \equiv \text{Prob}_G \left[t_k \varepsilon^l - t_k \varepsilon^j \leq x \quad \text{and} \quad l = \arg \max_i \left\{ u_k^i + t_k \varepsilon^i \right\} \quad \text{s.t.} \quad \hat{l} \neq j \right]$$

the probability (induced by the joint G) that the difference in consumers' brand tastes for firms l and j is less than $x \in \mathbb{R}$, and that firm l is j 's best competitor for type- k consumers. The demand from type- k consumers faced by firm j is then

$$D_k^j \left(u_k^j, \mathbf{u}_k^{-j} \right) = \sum_{l \neq j} H_k^{j,l} \left(u_k^j - u_k^l \mid \mathbf{u}_k^{-j} \right). \quad (2)$$

Our discrete-choice framework admits as special cases the Hotelling specification (where $J = 2$, $\varepsilon^1 \sim U[0, 1]$ and $\varepsilon^2 = 1 - \varepsilon^1$),¹⁵ the Probit specification (where ε^j 's are iid draws from a standard Normal cdf), and the Logit specification (where ε^j 's are iid draws from a standard Gumbel cdf). For instance, equation (2) simplifies to

$$D_k^j(u_k^j, u_k^l) = p_k \left(\frac{1}{2} + \frac{u_k^j - u_k^l}{2t_k} \right) \quad \text{and} \quad D_k^j(u_k^j, \mathbf{u}_k^{-j}) = \frac{p_k \exp\left(\frac{u_k^j}{t_k}\right)}{\sum_{l=1}^J \exp\left(\frac{u_k^l}{t_k}\right)}$$

in the Hotelling and Logit cases, respectively.¹⁶

A crucial feature of equation (2) is that firms' demands only depend on the (gross) indirect utilities offered by firms to each consumer type. Accordingly, we find it convenient to formulate the firms' problem in terms of indirect utilities, rather than price-quality pairs. To this end, the next lemma addresses the following question: which menu m maximizes a firm's profit conditional on delivering the indirect utility profile (u_l, u_h) ? Before stating the result, let us denote by

$$q_k^e \equiv \arg \max_q \theta_k q - \varphi(q) \quad \text{and} \quad S_k^e \equiv \theta_k q_k^e - \varphi(q_k^e)$$

the type- k 's efficient quality and efficient surplus.¹⁷ From now on, we drop the superscript j on firms' menus to lighten notation.

Lemma 0. [Incentive Compatibility] *Consider an equilibrium menu m , and let (u_l, u_h) be its profile of indirect utilities. Then the menu's qualities are given by*

$$q_l(u_l, u_h) = \begin{cases} \frac{u_h - u_l}{\Delta\theta} & \text{if } u_h - u_l < q_l^e \Delta\theta \\ q_l^e & \text{if } u_h - u_l \geq q_l^e \Delta\theta \end{cases} \quad \text{and} \quad q_h(u_l, u_h) = \begin{cases} \frac{u_h - u_l}{\Delta\theta} & \text{if } u_h - u_l > q_h^e \Delta\theta \\ q_h^e & \text{if } u_h - u_l \leq q_h^e \Delta\theta. \end{cases}$$

¹⁵More generally, the spokes model of Chen and Riordan (2007), which extends Hotelling to an arbitrary number of firms, is also a special case of our discrete-choice framework.

¹⁶In the Hotelling case, we set $D_k^j(u_k^j, u_k^l) = 0$ if $u_k^j - u_k^l < -t_k$ and $D_k^j(u_k^j, u_k^l) = 1$ if $u_k^j - u_k^l > t_k$.

¹⁷Because $\varphi'(0) = 0$ and $\lim_{q \rightarrow \infty} \varphi'(q) = \infty$, the efficient qualities and surpluses exist and are strictly positive for both consumer types.

Given (u_l, u_h) , one can determine via Lemma 0 the equilibrium quality levels (q_l, q_h) , and hence also the prices (y_l, y_h) of any equilibrium menu. It is therefore convenient to abuse notation and identify each menu to its indirect-utility profile: $m = (u_l, u_h)$. The surplus generated by each contract $k \in \{l, h\}$ in menu $m = (u_l, u_h)$ is then given by

$$S_k(u_l, u_h) \equiv q_k(u_l, u_h) \theta_k - \varphi(q_k(u_l, u_h)),$$

and the firm's profit per sale of this contract is $S_k(u_l, u_h) - u_k$.

For future comparison, it is useful to revisit the *monopolistic (or Mussa-Rosen)* outcome, formulated in the indirect-utility space. Namely, the monopolist menu, denoted by $m^\infty \equiv (u_l^\infty, u_h^\infty)$, solves

$$\max_{(u_l, u_h)} \{p_l(S_l(u_l, u_h) - u_l) + p_h(S_h(u_l, u_h) - u_h)\} \quad \text{s.t.} \quad u_h \geq u_l \geq 0. \quad (3)$$

As previous literature has shown, the solution to this problem entails $u_l^\infty = 0$ (low types obtain no rents) and $u_h^\infty < q_l^e \Delta \theta$ (constraint IC_h binds). Moreover, low types are not served ($q_l^\infty = 0$) in case $\theta_l - \frac{p_h}{p_l} \Delta \theta \leq 0$. Otherwise, $u_h^\infty > 0$ is implicitly given by

$$\frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, u_h^\infty) - 1 = 0, \quad (4)$$

in which case low-type consumers are offered a positive but inefficiently low quality level. By contrast, consumers appropriate the entire efficient surplus in the *perfectly competitive (or Bertrand) menu* $m^0 \equiv (S_l^e, S_h^e)$, where quality provision is efficient to consumers of all valuations and firms derive zero profits from each contract in the menu.

3.2 Balanced Oligopoly

Before analyzing the full-fledged model, we shall first revisit the case of a “balanced oligopoly,” where brand loyalty is invariant to type. To this end, we shall introduce the following regularity condition, which we retain for the rest of the paper.

Assumption 1. (log-concavity) Consider the cdf

$$H(x) \equiv \text{Prob}_G [\varepsilon^1 - \varepsilon^2 \leq x \mid \varepsilon^2 = \max\{\varepsilon^2, \dots, \varepsilon^J\}],$$

which density is $h(x)$. Then H is strictly log-concave.

Assumption 1 calls for the firms' demand to be strictly log-concave when all competitors play symmetric strategies. This requirement is satisfied by most specifications of interest (Hotelling Logit, Probit, etc) and is common in models of oligopolistic competition.¹⁸ For future use, let us denote by $\lambda \equiv \frac{H(0)}{h(0)}$ the *baseline markup*. To understand this terminology, note that, in the equilibrium under

¹⁸See, for instance, Anderson et al (1992).

complete information, all firms offer the efficient quality level and sell at a profit margin equal to λt_k . In the Hotelling specification, the baseline markup is $\lambda = 1$; in the Logit specification, it is $\lambda = \frac{J}{J-1}$.

The next proposition derives the equilibrium when brand preferences are equally intense across consumer types. For convenience, we let $\bar{\eta} \equiv S_h^e - q_l^e \Delta\theta$. This threshold measures the “slackness” of constraint IC_h , determining the smallest brand loyalty parameter that activates this constraint.

Proposition 0. (*Equilibrium: Balanced Oligopoly*) *Suppose the intensity of brand preferences is the same across types, and let $t \equiv t_l = t_h$. Then there exists a unique pure-strategy equilibrium, which is such that:*

(a) *If $t \in [0, \frac{\bar{\eta}}{\lambda}]$, $u_k^* = \max\{S_k^e - \lambda t, 0\}$ for $k \in \{l, h\}$. Quality provision is efficient.*

(b) *If $t > \frac{\bar{\eta}}{\lambda}$, $u_l^* = 0$ and, whenever positive, u_h^* is implicitly given by*

$$\frac{S_h^e - u_h^*}{\lambda t} + \frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, u_h^*) - 1 = 0. \quad (5)$$

Moreover, u_h^ is decreasing in t , and converges to the monopolistic level u_h^∞ as t grows unbounded. Quality is efficiently (resp., under-) provided to high-type (resp., low-type) consumers.*

When t is small (i.e., $t \leq \frac{S_l^e}{\lambda}$), the equilibrium is close to the perfectly competitive outcome, where neither individual rationality or incentive constraints bind. Accordingly, quality is efficiently provided to consumers of all valuations, and firms’ markups are constant across the product line (and equal to λt). This outcome coincides type-by-type with that under complete information. Moreover, it can be implemented by the “cost-plus-fee” tariff $T(q) \equiv \lambda t + \varphi(q)$, as first observed by Armstrong and Vickers (2001) and Roche and Stole (2002).

For t larger than $\frac{S_l^e}{\lambda}$, the individual rationality constraint is binding for low-valuation consumers, whose surplus is fully extracted by firms. As long as $t \leq \frac{\bar{\eta}}{\lambda}$, this is the only binding constraint, and quality provision remains efficient to both consumer types.¹⁹ Otherwise, the incentive constraint of high-valuation consumers also binds, and equilibrium is characterized by equation (5). As the intensity of brand preferences increases (i.e., t grows large), firms are able to extract more rents from high-valuation consumers, which requires decreasing the low-type quality away from its efficient level. In the limit as $t \rightarrow \infty$, equilibrium converges to the monopolistic outcome. This can be readily seen from equation (5): Its first term vanishes as t grows unbounded, making the equilibrium condition coincide with the monopolist’s optimality condition in equation (4).

Remark 1. (*Participation*) *It is common to assume that consumers have access to an outside option that delivers a certain level of utility (normalized to zero) against which they compare the*

¹⁹Indeed, $\bar{\eta} > S_l^e$, as implied by the convexity of the cost function φ . This implies that constraint IR_l binds “before” (i.e., for smaller t ’s) IC_h in the balanced-oligopoly case.

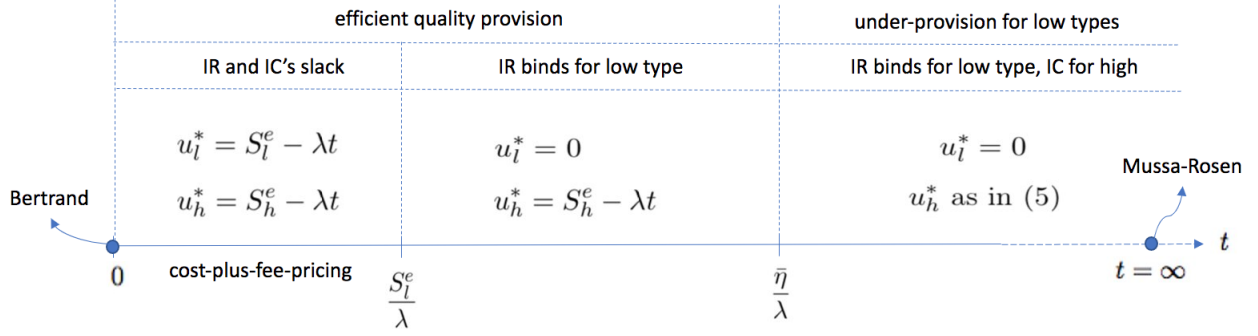


Figure 1: Equilibrium in a balanced oligopoly.

net utilities from each firm, as in (1) - call it the standard approach.²⁰ Instead, we model non-participation by endowing each firm with a no-purchase contract and requiring consumers to select one available contract (possibly a no-purchase one). This is the discrete-choice equivalent of the “co-located outside option” model proposed by Bénabou and Tirole (2016).²¹ This alternative approach is more suitable to study competition in the presence of price discrimination than the standard one, as it assures that changes in the intensity of brand preferences do not affect the consumers’ decision to withdraw from the market. To better understand this point, note that the standard approach implies that those consumers who are indifferent between two goods (say, a GM and a Ford car) are the ones who most value the outside option (say, public transport), and the more so the larger the parameter t is. Accordingly, in the context of Proposition 0, when t is slightly above $\frac{2}{3}S_l^e$, firms are local monopolists for low-type consumers (as those consumers with x close to $\frac{1}{2}$ left the market), but compete under full market coverage for high types. As a result, variations in t affect the intensity of brand preferences of the latter (whose relevant comparison is across firms’ contracts), while affecting consumer participation for the former (whose relevant comparison is between the closest firm and the outside option). As t grows large, the volume of sales to consumers (of any type) shrinks to zero. In the alternative approach followed here, by contrast, the value of the no-purchase option (relative to the closest brand) is not affected by changes in t , which therefore can be identified with the intensity of brand preferences. As a result, firms are always in competition for both consumer types. In contrast to previous work, our model covers the whole spectrum of competitive intensity, converging to the Mussa-Rosen outcome as consumers become intensely brand loyal.

²⁰This is the case of Roche and Stole (2002) and Armstrong and Vickers (2001) as well as of most related literature, as discussed in Section 7.

²¹Bénabou and Tirole provide an spatial interpretation according to which each consumer has to travel to a shopping mall (there are two, each located in one extreme of the Hotelling segment) to purchase either a product of one of the competing firms or some unmodelled outside good. For a probabilistic interpretation, see Poletti and Wright (2004), who propose an equivalent formulation.




Figure 2: Equilibrium distortions and brand substitutability across types.

4 Equilibrium

Just like in some markets low-valuation consumers have stronger brand preferences than those with high-valuations (e.g., cable TV), in others the reverse pattern is verified (e.g., air travel). We now study equilibrium outcomes in the full-fledged model, where brand loyalty is positive ($t_l, t_h > 0$) and type-dependent. In this section, we focus on pure-strategy equilibria, whenever such equilibria exist.

4.1 Characterization

To describe equilibrium patterns, consider the thresholds $\eta_h \equiv S_h^e - q_h^e \Delta\theta > 0$ and $\eta_l \equiv S_l^e - q_l^e \Delta\theta$, which measure the “slackness” of constraint IC_l and IC_h , respectively, when brand loyalty differs across types. To streamline the exposition, we assume that:

Assumption 2. $\eta_h > 0$.

When the cost function has the power form, $\varphi(q) = \frac{1}{a}q^a$, where $a > 1$, Assumption 2 is satisfied if and only if $a\theta_l > \theta_h$. Intuitively, this assumption requires consumer types to be sufficiently close so that the self-selection constraint IC_l may bind in equilibrium. If this condition is violated, we obtain the less interesting case where the constraint IC_l is slack for all brand loyalty profiles (t_l, t_h) .

It is useful to define the function $\Lambda(t_l, t_h) \equiv \lambda t_h + \max\{S_l^e - \lambda t_l, 0\}$. After noting that $\eta_h < \bar{\eta}$,²² consider the following parameter regions, illustrated in Figure 2:

$$E \equiv \{(t_l, t_h) \in \mathbb{R}_{++}^2 : \eta_h \leq \Lambda(t_l, t_h) \leq \bar{\eta}\},$$

$$D_+ \equiv \{(t_l, t_h) \in \mathbb{R}_{++}^2 : \Lambda(t_l, t_h) < \eta_h\}, \quad \text{and} \quad D_- \equiv \{(t_l, t_h) \in \mathbb{R}_{++}^2 : \Lambda(t_l, t_h) > \bar{\eta}\}.$$

The next proposition characterizes pure-strategy equilibria in each of these regions.

²²Indeed, by definition, $\eta_h = S_h^e - q_h^e \Delta\theta < S_h^e - q_l^e \Delta\theta = \bar{\eta}$.

Proposition 1. (Characterization) Suppose $(t_l, t_h) \in \mathbb{R}_{++}^2$. Then:

(a) *Region E:* There is a unique pure-strategy equilibrium, in which $u_k^* = \max\{S_k^e - \lambda t_k, 0\}$ for each $k \in \{l, h\}$. Quality provision is efficient to both consumer types, who obtain the same payoffs as under complete information.

(b) *Region D_+ :* A pure-strategy equilibrium exists if $t_h \geq \tau_h(t_l)$, where the threshold $\tau_h(t_l)$ is such that $(t_l, \tau_h(t_l)) \in D_+$. This equilibrium is unique and (u_l^*, u_h^*) solve

$$S_h(u_l^*, u_h^*) - u_h^* + \lambda t_h \left(\frac{\partial S_h}{\partial u_h}(u_l^*, u_h^*) - 1 \right) = 0 \quad \text{and} \quad \frac{S_l^e - u_l^*}{\lambda t_l} + \left(\frac{p_h}{p_l} \frac{\partial S_h}{\partial u_l}(u_l^*, u_h^*) - 1 \right) \leq 0,$$

where the second condition is an (in)equality if $u_l^* > 0$ ($u_l^* = 0$). Quality is over-provided to high types, but efficiently provided to low types. Relative to the complete information benchmark, high types lose, while low types gain from asymmetric information.

(c) *Region D_- :* A pure-strategy equilibrium exists if $t_l \geq \tau_l(t_h)$, where the threshold $\tau_l(t_h)$ is such that $(\tau_l(t_h), t_h) \in D_-$. This equilibrium is unique and (u_l^*, u_h^*) solve

$$\frac{S_h^e - u_h^*}{\lambda t_h} + \left(\frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(u_l^*, u_h^*) - 1 \right) \leq 0 \quad \text{and} \quad S_l(u_l^*, u_h^*) - u_l^* + \lambda t_l \left(\frac{\partial S_l}{\partial u_l}(u_l^*, u_h^*) - 1 \right) \leq 0,$$

where the first condition is an (in)equality if $u_h^* > 0$ ($u_h^* = 0$), while the second condition is an (in)equality if $u_l^* > 0$ ($u_l^* = 0$). Quality is under-provided to low types, but efficiently provided to high types. Relative to the complete information benchmark, high types gain, while low types lose from asymmetric information.

Recall from Proposition 0(a) that IC's are slack in the “diagonal” (i.e., when $t_l = t_h$) provided the intensity of brand preferences is low (in the sense that $t_l = t_h < \frac{S_l^e}{\lambda}$). The equilibrium is then identical to that under complete information. Proposition 1(a) shows that the complete-information outcome is an equilibrium provided brand preferences are not “too different” across consumer types (in which case IC's remain slack). The same holds true over the horizontal band in region E (in which $t_l \geq \frac{S_l^e}{\lambda}$). The reason is that raising t_l , while keeping t_h constant, does not affect firm's incentives (as u_l^* is already at zero). So the balanced-oligopoly equilibrium remains (the unique) equilibrium, and quality is efficiently provided to both consumer types.

Consider now region D_+ . There, high types are sufficiently less brand-loyal than low types, who constitute the market segment with higher potential for profits. To prevent profit dissipation, whereby low types migrate to the high-type contract, firms set the quality of the “premium” product inefficiently high, which renders this product less attractive to low types. In the simpler case where

$u_l^* = 0$, the equilibrium condition from Claim (b) reveals that u_h^* is determined by

$$\underbrace{\frac{S_h(0, u_h^*) - u_h^*}{\lambda t_h}}_{\text{poaching gain}} - \underbrace{1}_{\substack{\text{mark-up} \\ \text{loss per sale}}} + \underbrace{\left(\frac{\partial S_h}{\partial u_h}(0, u_h^*)\right)}_{\text{efficiency loss per sale}} = 0. \quad (6)$$

Intuitively, when choosing how much utility to leave to high types, firms balance the gains from poaching consumers away from the competitor, which is the first term in (6), with the per sale loss from reducing the price, which is the second term, compounded with the efficiency loss from tightening the incentive constraint, which is the last term. The latter is absent in the complete information benchmark, what explains why informational rents are negative for high-type consumers:

$$u_h^* = S_h(u_l^*, u_h^*) - \lambda t_h + \lambda t_h \frac{\partial S_h}{\partial u_h}(u_l^*, u_h^*) \leq S_h^e - \lambda t_h.$$

Intuitively, firms have less incentives to increase the high-type payoff relative to the complete information benchmark. The reason is that IC_l binds, so raising u_h tightens this constraint, thus decreasing the efficiency of the high-type contract. The opposite applies to low types, for which firms are more compelled to provide rents (so as to relax this constraint).

By contrast, in region D_- , low types are significantly less brand-loyal than high types, who constitute the market segment with higher profit potential. Accordingly, firms under-provide quality at the bottom of the product line, which prevents high types from purchasing the low-quality good (which exhibits a smaller profit margin).

The pattern of informational rents is the mirror image of that obtained in region D_+ . Namely, relative to the complete information benchmark, high types gain, while low types lose from asymmetric information. Intuitively, firms have an extra incentive to increase high-type payoffs relative to the complete information benchmark. On top of the usual poaching gains and mark-up losses, increasing u_h relaxes the binding constraint IC_h , thus increasing the efficiency of the low-type contract. The opposite applies to low types, for which firms are less compelled to provide rents (so as not to tighten this constraint).

Interestingly, the latter effect is shrouded in the balanced oligopoly case, where IC_h binds only when $u_l^* = 0$ (as, note from Figure 2, $t_l > S_l^e$ whenever the 45-degree line belongs to region D_-). In this case, low types obtain the same payoff as under complete information (zero), while high-types obtain positive informational rents.

It is worth noting that Proposition 1 guarantees the existence of a pure-strategy equilibrium in region E , and in regions D_+ and D_- provided one is “close” to region E . It turns out that pure-strategy equilibria fail to exist if brand loyalty is sufficiently different across types (i.e., fixing $t_l > 0$, t_h is sufficiently small, and vice-versa). We will come back to this point in next section, clarifying where and why non-existence obtains. Before doing so, the next subsection leverages on

the characterization of Proposition 1 to understand how changes in brand loyalty affect equilibrium outcomes.

4.2 Comparative Statics

Proposition 2. (Comparative Statics on Quality and Payoffs) *Consider a neighborhood around $(t_l, t_h) \in \mathbb{R}_{++}^2$ where the pure-strategy equilibrium exists. Then, for $k \in \{l, h\}$,*

$$\frac{\partial u_k^*}{\partial t_h}, \frac{\partial u_k^*}{\partial t_l} \leq 0, \quad \text{and} \quad \frac{\partial q_k^*}{\partial t_h} \leq 0 \leq \frac{\partial q_k^*}{\partial t_l}.$$

If $(t_l, t_h) \notin E$ and $u_l^ > 0$, then $\frac{\partial u_k^*}{\partial t_h}, \frac{\partial u_k^*}{\partial t_l} < 0$ for $k \in \{l, h\}$. In this case, if $q_k^* \neq q_k^e$ then $\frac{\partial q_k^*}{\partial t_h} < 0 < \frac{\partial q_k^*}{\partial t_l}$.*

Not surprisingly, the equilibrium indirect utility of both consumer types strictly decreases as brand preferences (of either type) become more intense. More interesting, perhaps, is the effect on qualities when some incentive constraint binds (otherwise, quality levels are efficient). In this case, as high types develop more intense brand preferences, equilibrium quality levels go down. The reason is the following: changes in t_h directly affect competition for high types, but only indirectly for low types (through incentive constraints). As a result, an increase in t_h decreases u_h^* *faster* than u_l^* , what implies that the quality of the inefficient contract decreases. When constraint IC_l is binding (as in region D_+), it is the high-type quality that is inefficient. As such, q_h^* goes down (reducing the distortion) as t_h goes up, whereas q_l^* remains constant at its efficient level. In turn, when constraint IC_h is binding (as in region D_-), it is the low-type quality that is inefficient. As such, q_l^* goes down (magnifying the distortion) as t_h goes up, whereas q_h^* remains constant at its efficient level. Therefore, variations in t_h can either increase or decrease equilibrium welfare, depending on whether the preference profile (t_l, t_h) belongs to regions D_+ or D_- . When (t_l, t_h) lies in the efficient region E , variations in t_h have no effect on equilibrium qualities.

Mutatis mutandis, the same logic explains the effect of t_l on equilibrium quality levels. Because an increase in t_l decreases u_l^* *faster* than u_h^* , the quality of the inefficient contract (if positive) shall increase. Accordingly, an increase in t_l strictly increases q_h^* when $(t_l, t_h) \in D_+$ (magnifying the distortion), but strictly increases q_l^* when $(t_l, t_h) \in D_-$ and $q_l^* > 0$ (reducing the distortion).

Accordingly, Proposition 2 reveals that, in the presence of self-selection constraints, competition and welfare are often misaligned, in that more competitive markets (in the sense that consumers are less brand-loyal) often produce lower welfare. At the heart of the matter lies the idea that, under asymmetric information, contract offers are interdependent across consumer segments. This interdependency renders competition welfare-decreasing whenever it tightens incentive constraints.

In light of this discussion, it is natural to inquire how a merger, creating a monopoly that designs and prices both firms' product lines, affects consumer surplus and welfare. To this end, note that the profit-maximization problem of the merged entity (controlling firms a and b) coincides with that

of a Mussa-Rosen monopolist, described in equation (3).²³ Conveniently, the Mussa-Rosen outcome (applied to each firm) is the limit equilibrium as the brand loyalty parameters t_l and t_h grow large (as shown in Proposition 0). It then follows that we can perform merger analysis by comparing the equilibrium outcome at a given profile (t_l, t_h) to the limit outcome as $t_l, t_h \rightarrow \infty$.

The comparative statics developed in Proposition 2 come in handy for this purpose. In particular, it implies that the merger always decreases consumer surplus, as the equilibrium payoffs u_l^* and u_h^* are decreasing in the brand loyalty parameters. The effect on welfare is more nuanced, as explored in the next corollary.

Corollary 1. (*Welfare Effect of a Merger*) Consider a brand loyalty profile $(t_l, t_h) \in \mathbb{R}_{++}^2$ where a pure-strategy equilibrium exists. If $(t_l, t_h) \in D_-$, the merger from oligopoly to monopoly decreases welfare. If, however, $(t_l, t_h) \in D_+$, the merger may increase welfare provided low types are sufficiently brand-loyal, while high types are little brand-loyal.

According to Corollary 1, if the oligopolistic outcome involves downward distortions at the bottom of the product line, the merger decreases welfare. The reason is that, under oligopoly, the quality of the baseline good is always greater than under monopoly. Intuitively, business stealing creates an extra incentive (relative to monopoly) for the competing firms to relinquish indirect utility to high types, which relaxes the incentive constraint.

The merger can however increase welfare if the oligopoly outcome exhibits over-provision at the top of the product line. To understand why, note that, in this case, the merger substitutes one distortion (under-provision at the bottom) for another (over-provision at the top). As implied by Proposition 2, the latter distortion is more detrimental to welfare when low (resp., high) types are sufficiently (resp., little) brand loyal. In this case, the welfare comparison favors the monopoly outcome (for a wide range of parameters).²⁴

The comparative statics on prices is explored in the following result.

Proposition 3. (*Comparative Statics on Prices*) Consider a neighborhood around $(t_l, t_h) \in \mathbb{R}_{++}^2$ where the pure-strategy equilibrium exists and $q_l^* > 0$, and denote by (y_l^*, y_h^*) the equilibrium price profile.

(a) If $(t_l, t_h) \in E$, then $y_k^* = \varphi(q_k^e) + \min\{\lambda t_k, S_k^e\}$.

(b) If $(t_l, t_h) \in D_+$, then

$$\frac{\partial y_l^*}{\partial t_l}, \frac{\partial y_l^*}{\partial t_h}, \frac{\partial y_h^*}{\partial t_l} \geq 0,$$

with strict inequality if and only if the constraint IR_l is slack. Moreover, y_h^* is decreasing in t_h if IR_l binds, but is quasi-convex in t_h if IR_l is slack and $\varphi'''(q) \leq 0$.

²³This follows from noting that the merged entity optimally chooses the same menu for firms a and b .

²⁴The exact condition, which depends on the support of types $\{\theta_l, \theta_h\}$ and their respective probabilities, appears in the Online Appendix.

(c) If $(t_l, t_h) \in D_-$, then

$$\frac{\partial y_h^*}{\partial t_h} > 0, \quad \text{and} \quad \frac{\partial y_l^*}{\partial t_l}, \frac{\partial y_h^*}{\partial t_l} \geq 0,$$

with strict inequality if and only if the constraint IR_l is slack. Moreover, y_l^* is decreasing in t_h if IR_l binds, but is quasi-convex in t_h if IR_l is slack and $\varphi'''(q) \leq 0$.

Prices always increase with the intensity of brand preferences when the quality of the product is set efficiently. This familiar intuition explains why prices increase with (t_l, t_h) in region E , and why the baseline price y_l^* (resp., premium price y_h^*) increase with (t_l, t_h) in region D_+ (resp., D_-).

The analysis is more subtle when changes in the intensity of brand preferences jointly affect quality provision and utility levels. This occurs, for instance, with the premium product when $(t_l, t_h) \in D_+$. As t_l increases, high types are worse-off (u_h^* decreases), whereas the quality of the premium product increases (q_h^* increases), so y_h^* has to increase as well. By contrast, as t_h increases, high-type payoffs decrease, as so does the quality of the premium product. The latter effect dominates when IR_l binds, which implies the premium product becomes cheaper as high types become more brand loyal. When IR_l is slack, this pattern is more nuanced, as the premium price y_h^* is U-shaped in the brand loyalty parameter t_h . To prove this global quasi-convexity result, we need to impose a weak regularity condition (namely, that $\varphi'''(q) \leq 0$).

A similar logic explains why, in region D_- , the baseline product may become cheaper as high types becomes more brand-loyal. When IR_l binds, the baseline price goes down with t_h because the quality of the baseline product falls more than the equilibrium payoff of low types. When IR_l is slack, the race between quality and payoff changes leads to U-shaped pattern.

More broadly, Proposition 3 reveals that, under self-selection constraints, variations in the level of prices are a misleading indicator of the degree of competition in the market. This is consistent with the ambiguous relationship found in the empirical literature between the degree of competition and the level of prices in markets characterized by self-selection. For instance, Chu (2010) documents that cable companies in the US reacted to new competition by satellite television by raising both price and quality (as determined by the number of available channels), with consumers benefiting overall from the higher-priced offerings.

5 Dispersion of Offers

The previous section studied pure-strategy equilibria, revealing their properties and performing comparative statics. Yet, such equilibria may fail to exist, in which case we have to look for mixed-strategy equilibria, where firms randomize over menus of contracts. In this section, we start by investigating the reasons for non-existence of pure-strategy equilibria, to then study equilibria under mixed strategies.

To understand the non-existence phenomenon, it is instructive to look at the following special case of our model, where there is perfect competition for one consumer type.

5.1 Bottom-of-Barrel Oligopoly

A bottom-of-barrel oligopoly bears its name due to the fact that brand preferences are stronger among those consumers who have the lowest willingness to pay for quality (and therefore the lowest potential for profits). To capture this possibility in the starkest manner, we assume that high-valuation consumers see firms as perfect substitutes, i.e., $t_h = 0$. In turn, the intensity of brand preferences among low-valuation consumers is unconstrained, as t_l is allowed to take any positive value.

Before investigating equilibrium outcomes, it is convenient to define the *zero-profit h-type utility* \hat{u}_h as the implicit solution to

$$S_h(0, \hat{u}_h) - \hat{u}_h = 0. \quad (7)$$

In words, \hat{u}_h is the highest indirect utility that firms can relinquish to high-type consumers while obtaining zero profit from the high-type contract and fully extracting rents from low types. By Assumption 2, $\hat{u}_h > q_h^e \Delta\theta$, which implies that the high-type quality of the menu $(0, \hat{u}_h)$ is distorted upwards: $q_h(0, \hat{u}_h) > q_h^e$.

The next proposition clarifies when a pure-strategy equilibrium exists.

Proposition 4. (*Pure-Strategy Equilibrium: Existence*) *Suppose there is perfect competition for high types ($t_h = 0$), but imperfect for low types ($t_l > 0$). Then:*

(a) *If $0 \leq t_l \leq \frac{S_l^e - \eta_h}{\lambda}$ there is a unique pure-strategy equilibrium, with $u_l^* = S_l^e - \lambda t_l$ and $u_h^* = S_h^e$. Quality provision is efficient.*

(b) *No pure-strategy equilibrium exists if $\frac{S_l^e - \eta_h}{\lambda} < t_l < \tilde{t}_l$, where the threshold:*

$$\tilde{t}_l \equiv \inf \left\{ t_l : \frac{\partial S_h}{\partial u_l}(0, \hat{u}_h) < \frac{1}{J} \frac{p_l}{p_h} \left(1 - \frac{S_l^e}{\lambda t_l} \right) \right\},$$

and $\tilde{t}_l \equiv \infty$ if the inequality inside brackets is violated for all $t_l > 0$.

(c) *If $t_l \geq \tilde{t}_l$, there is a unique pure-strategy equilibrium, with $u_l^* = 0$ and $u_h^* = \hat{u}_h$. Quality is over-provided to high types, but efficiently provided to low types.*

As illustrated in Figure 3, Proposition 4 identifies three regions. When t_l is small, we are in region *E* of Proposition 1. Accordingly, an equilibrium in pure strategies exists, which is unique and coincides with that under complete information. Because $t_l > t_h = 0$, firms obtain zero profits from high types, but a positive profit from low types.

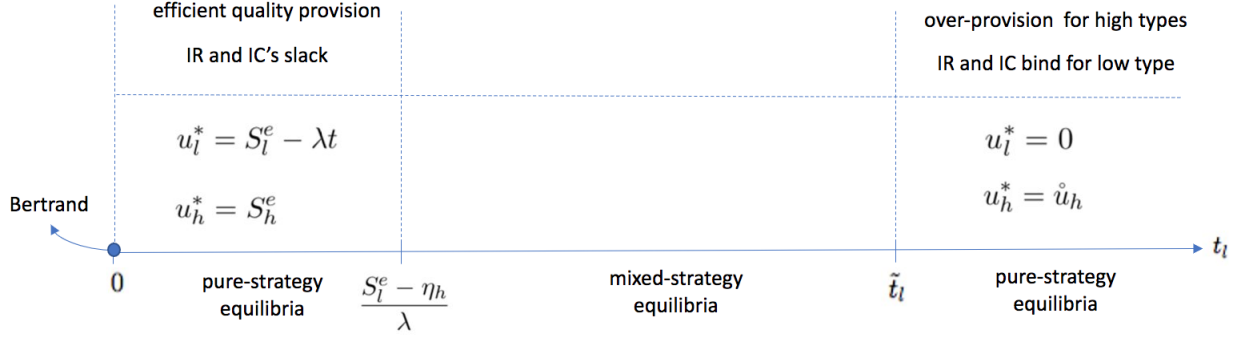


Figure 3: Equilibrium in a bottom-of-barrel oligopoly.

When t_l exceeds the threshold $\frac{S_l^e - \eta_h}{\lambda}$, the “complete-information” equilibrium described above can no longer be sustained, as constraint IC_l would be violated. This means that no pure-strategy equilibrium exhibits efficient quality provision. Perfect competition, however, implies that firms obtain zero profit from high-type consumers. As a consequence, firms can ignore the existence of high-type consumers and respond *as if* they were playing a competition game where only low types are present. These two observations imply that, if an equilibrium in pure strategies exists, it has to satisfy

$$u_l^* = S_l^e - \lambda t_l \quad \text{and} \quad S_h(u_l^*, u_h^*) - u_h^* = 0. \quad (8)$$

Crucially, Assumption 2 implies that the high-type quality is above its efficient level in this putative equilibrium.

There is always a profitable deviation from this putative equilibrium provided $\frac{S_l^e - \eta_h}{\lambda} < t_l < \tilde{t}_l$. It works as follows: the deviating firm grants a small discount $\delta > 0$ to low types, which relaxes the IC_l constraint. This enables the firm to reduce the quality provided to high types, therefore increasing the efficiency from their respective contracts (recall there was over-provision in the putative equilibrium). Because there is perfect competition among high types, the deviating firm can then adjust prices to slightly undercut its rival, conquering the whole high-type market and appropriating the correspondent efficiency gain. That this deviation is profitable comes from the fact that the profit gain among high types is of first-order magnitude, while the discount δ to low types entails only a second-order profit loss.²⁵ We refer to this strategy as the *relax-and-undercut* deviation, as it involves relaxing incentive compatibility to enable undercutting the rival firm.

For $t_l > \frac{S_l^e}{\lambda}$, the IR_l constraint necessarily binds. By the same reasoning above, if an equilibrium in pure strategies exists, it has to be such that $(u_l^*, u_h^*) = (0, \hat{u}_h)$, where \hat{u}_h is zero-profit h -type utility given by equation (7). In this menu, because IC_l binds, high types are provided inefficiently high quality, while appropriating the full (inefficient) surplus produced by their contract. Low types

²⁵To see why, note that $u_l^* = S_l^e - \lambda t_l$ in an interior optimum, therefore being a local maximand.

endure full rent extraction, as $u_l^* = 0$, and are offered their efficient quality level.

The relax-and-undercut deviation now produces a first-order profit gain among high-types at the expense of a *first-order profit loss* among low types. The loss is now first-order because the putative equilibrium utility $u_l^* = 0$ is at the corner dictated by the IR_l constraint (therefore exhibiting a non-zero shadow cost). The race between these two effects is resolved in favor of deviating if and only if t_l is below the threshold \tilde{t}_l . Intuitively, if brand preferences are mild (in the sense that $t_l < \tilde{t}_l$), the business-stealing effect from discounting the low-type price is sufficiently large to render the relax-and-undercut deviation profitable. This is always the case when $\tilde{t}_l = \infty$.²⁶

In sum, Proposition 4 reveals that, whenever IC_l is the only constraint binding, a pure-strategy equilibrium fails to exist. Crucially, non-existence is solely due to the interplay between heterogenous brand loyalties and self-selection by consumers. In this case, the equilibrium exhibits dispersion of offers, whereby firms randomize their choice of menus.

Remark 2. (Cream Skimming) *Mirror image of the bottom-of-barrel, the cream-skimming oligopoly exhibits perfect competition for low types, $t_l = 0$, but imperfect competition for high types, $t_h > 0$. Accordingly, firms' market power is stronger among those consumers who have the highest willingness to pay, exhibiting the largest potential for profits. Similarly to Proposition 4, no pure-strategy equilibrium exists if $\frac{S_h^e - \eta}{\lambda} < t_h < \tilde{t}_h$, where the threshold \tilde{t}_h satisfies*

$$\tilde{t}_h \equiv \inf \left\{ t_h : \frac{\partial S_l}{\partial u_h}(0, 0) < \frac{1}{J} \frac{p_h}{p_l} \left(1 - \frac{S_h^*}{\lambda t_h} \right) \right\}.$$

An analogous relax-and-undercut deviation improves upon the putative equilibrium outcome, where firms obtain zero profit from low types and compete for high types as under complete information. This deviation consists on granting a small discount to high-type consumers, which relaxes the IC_h constraint and enables the deviating firm to increase the quality provided to low types (which is underprovided in the putative equilibrium). Because there is perfect competition for these consumers, the deviating firm can then adjust prices to slightly undercut its rival, conquering the whole low-type market and appropriating the correspondent efficiency gain. As in the bottom-of-barrel case, this deviation trades off a first-order efficiency gain among low types with a profit loss from high types (which is second-order if IR_h is slack). Pure-strategy equilibria may be restored only if IR_h binds, as captured by the threshold \tilde{t}_h .

Remark 3. (Non-Existence: General Case) *The non-existence argument described above extends beyond the bottom-barrel and cream-skimming cases. On a more technical level, the non-existence of a pure-strategy equilibrium stems from the fact that firms' best responses are not quasi-concave. For instance, consider the region D_+ and fix some t_l satisfying $\frac{S_l^e - \eta_h}{\lambda} < t_l < \tilde{t}_l$. As shown*

²⁶Whether \tilde{t}_l is finite or not depends on parameters. For instance, if the cost is quadratic, $\varphi(q) = \frac{1}{2}q^2$, $\tilde{t}_l = \infty$ if $\kappa_l \equiv 1 - J \frac{p_h}{p_l} \left(\frac{2\theta_l - \theta_h}{\Delta\theta} \right) \leq 0$, but equals $\tilde{t}_l = \frac{\theta_l}{\lambda\kappa_l} < \infty$ if $\kappa_l > 0$.

in a previous version of this paper,²⁷ there exists a threshold $\hat{\tau}_h(t_l) > 0$ such that no pure-strategy equilibrium exists provided $t_h < \hat{\tau}_h(t_l)$. Indeed, for t_h low enough, firms' best responses fail to be locally quasi-concave at the putative equilibrium of Proposition 1. In this case, the putative equilibrium is a saddle point of the best response, thus exhibiting a local profitable deviation. Only for t_h large enough the global quasi-concavity of best-responses is restored, which guarantees that a pure-strategy equilibrium exists.²⁸ *Mutatis mutandis*, a similar pattern is found in region D_- .

5.2 Mixed-Strategy Equilibria

We now study mixed-strategy equilibria. As the next proposition reveals, any such equilibrium exhibits the following property: across any two menus offered in equilibrium, the indirect utilities offered to low- and high-type consumers co-move. This is the subject of the next definition, first proposed by Garrett et al (2019):

Definition 1. [Ordered Equilibrium] *A mixed-strategy equilibrium is said to be ordered if, for any two menus $\mathcal{M} = (u_l, u_h)$ and $\mathcal{M}' = (u'_l, u'_h)$ offered in equilibrium, $u_l < u'_l$ if and only if $u_h < u'_h$. In this case, the menu (u'_l, u'_h) is said to be more generous than the menu (u_l, u_h) .*

Accordingly, in an ordered mixed-strategy equilibrium, firms randomize over the generosity of the menu, which can be identified with the indirect utility relinquished to either low or high types (as these co-move across menus). The next proposition clarifies when mixed-strategy equilibria exist, and reveals that the pattern of distortions is parallel to that found in pure-strategy equilibria.

The result however requires two regularity conditions. The first strengthens the log-concavity requirement of Assumption 1.

Assumption 3. (generalized log-concavity) *Whenever $D_k^j(w_k^j, \mathbf{u}_k^{-j}) > 0$, the demand semi-elasticity $\frac{\partial \log D_k^j}{\partial u_k^j}(w_k^j, \mathbf{u}_k^{-j})$ is increasing in \mathbf{u}_k^{-j} .*

For two firms, Assumption 3 is implied by Assumption 1, but it is a little stronger when $J > 2$. Assumption 3 is satisfied by all specification of interest (Hotelling, Spokes, Logit, Probit, etc).

The second regularity condition requires that the complete information payoff of each firm j is quasi-concave in w_k^j for all (possibly mixed) strategies followed by the competing firms.

Assumption 4. (quasi-concavity) *For any distribution F_k^{-j} over \mathbf{u}_k^{-j} , the complete-information payoff*

$$\Phi(w_k^j) \equiv \mathbb{E}_{F_k^{-j}} \left[D_k^j(w_k^j, \mathbf{u}_k^{-j}) (S_k^e - w_k^j) \right]$$

is strictly quasi-concave in w_k^j .

²⁷See Gomes, Lozachmeur and Maestri (2020).

²⁸Of course, pure-strategy equilibria may exist if best responses fail global quasi-concavity. In this case, one has to compare the putative equilibrium profit with that obtained by the best non-local deviation. As shown in a previous version of this paper, under natural assumptions, the best non-local deviation is the relax-and-undercut strategy.

It is easy to verify analytically that Assumption 4 holds true in the case of Hotelling and Spokes.²⁹ We can then state the main result of this section.

Proposition 5. (*Mixed Strategies*) *Suppose $(t_l, t_h) \in \mathbb{R}_{++}^2$. An equilibrium (either in pure or mixed strategies) exists, and every mixed-strategy equilibrium is ordered.*

- (a) *Region E: no mixed strategy-equilibria exist.*
- (b) *Region D_+ : A mixed-strategy equilibrium exists if $t_h < \hat{\tau}_h(t_l)$, where the threshold $\hat{\tau}_h(t_l) > 0$ whenever $t_l \in (S_l^e - \eta_h, \tilde{t}_l)$. In all equilibrium menus, quality is over-provided to high types, but efficiently provided to low types.*
- (c) *Region D_- : A mixed-strategy equilibrium exists if $t_l < \hat{\tau}_l(t_h)$, where the threshold $\hat{\tau}_l(t_h) > 0$ whenever $t_h \in (S_h^e - \eta_l, \tilde{t}_h)$. In all equilibrium menus, quality is under-provided to low types, but efficiently provided to high types.*

As revealed by Proposition 5, mixed-strategy equilibria are found in a neighborhood of any brand loyalty profiles $(t_l, 0)$ and $(t_h, 0)$ (corresponding to the bottom-barrel and cream-skimming cases) for which no pure-strategy equilibrium exists. To understand the ordered nature of mixed-strategy equilibria, consider a profile $(t_l, t_h) \in D_+$ in which such an equilibrium exists. As firms' market power is higher among low-type consumers, the constraint IC_l binds in all equilibrium menus, which therefore exhibit over-provision of quality in the high-type contract. This observation is key to understand the ordered property of equilibrium, which is intimately related to the relax-and-undercut deviation described above. Intuitively, firms differentiate themselves according to how big is the “discount” they give to low-type consumers. Those firms who grant the largest discounts are able to undertake the greatest reductions in the quality provided to high types, therefore obtaining the largest welfare gains. The larger is the surplus produced by the h -type contract, the larger is the incentive to relinquish more indirect utility to high types, so as to expand demand. As a consequence, the firms who offer the highest indirect utilities to low types (i.e., those who “discount” more) are also the ones that offer the highest indirect utilities to high types, i.e., menus are ordered. In equilibrium, all firms similarly randomize over a collection of ordered menus that maximize expected profits.

Proposition 5 accords with the fast-moving nature of many markets where product features and prices are typically “unstable” and short-lasting. For these markets, our model produces two empirical implications. First, dispersion of offers requires that different types exhibit (sufficiently) different propensities to switch brands. Second, whenever dispersion occurs, the gross utilities offered by firms are similarly ranked across the product line (as follows from the ordered property of equilibria).

²⁹For the Logit and Probit specifications, we have to proceed numerically. Our computational procedure suggests Assumption 4 also holds in these cases.

6 Discussion

6.1 Continuum of Types

The equilibrium characterization described above proceeded under the assumption of two consumer types, with no restrictions on the probability distribution. Studying models with more than two types is conceptually straightforward, but often challenging analytically. For instance, if in equilibrium only adjacent incentive constraints bind, the characterization of Proposition 1 can be easily extended to an arbitrary number of discrete types. If, however, consumer brand loyalty varies non-monotonically with type, the equilibrium is likely to exhibit bunching, which prevents a closed-form characterization. Crucially, when brand loyalty is type-dependent, it is elusive to find a condition on primitives that rules out bunching in equilibrium.

When types belong to a continuum, the equilibrium, rather than being described by a non-linear system of equations (as in Proposition 1), is characterized by a differential equation. This makes it easy verifying numerically that the equilibrium outcome exhibits no bunching. We can then explore the differential equation to confirm the main findings (regarding distortions and informational rents) of the binary-type model. This subsection illustrates this point assuming that types are uniformly distributed over some interval $[\underline{\theta}, \bar{\theta}] \subseteq \mathbb{R}_+$ and the cost function is quadratic: $\varphi(q) = \frac{1}{2}q^2$. It also employs numerical techniques to show that the non-existence phenomenon discussed in Section 5 extends to the continuum-type setting.

To facilitate comparison with the previous literature, we consider the Hotelling setup described in subsection 3.1. As before, we let the intensity of brand preferences change with one's preferences for quality, as described by the brand loyalty schedule $t(\theta)$. For simplicity, we assume that $t(\theta)$ is affine:

$$t(\theta) = \underline{t} + (\bar{t} - \underline{t}) \left(\frac{\theta - \underline{\theta}}{\bar{\theta} - \underline{\theta}} \right), \quad (9)$$

parametrized by $\underline{t}, \bar{t} \geq 0$. Note that $t(\bar{\theta}) = \bar{t}$ and $t(\underline{\theta}) = \underline{t}$, which explains the notation. If $\Delta t \equiv \bar{t} - \underline{t} < 0$, brand loyalty is decreasing in preferences for quality, and increasing if $\Delta t > 0$.

Roche and Stole (2002) showed that, when $\Delta t = 0$ and $t \equiv \bar{t} = \underline{t}$ is sufficiently small, the equilibrium is in pure strategies, and quality provision is efficient to all types ($q^*(\theta) = \theta$). Allowing for correlation between brand loyalty and brand preferences, our model also admits a pure-strategy equilibrium (assuming $|\Delta t| \neq 0$ is small). Moreover, when correlation is positive, i.e., $\Delta t > 0$ (resp., negative, i.e., $\Delta t < 0$), almost every quality is distorted downwards (resp., upwards), which is in line with Proposition 1. Figure 4 numerically illustrates this finding for $[\underline{\theta}, \bar{\theta}] = [1, 2]$.

Furthermore, when $\Delta t > 0$ (resp., $\Delta t < 0$) low types are worse-off (resp., better-off) under asymmetric information, while high types are better-off (resp., worse-off). Intuitively, when $\Delta t > 0$, competition for low types is hindered by the fact that high types have strong brand loyalty. To prevent profit dissipation (due to high types selecting low-quality contracts), firms then provide less




Figure 4: Equilibrium quality schedules for $[\underline{\theta}, \bar{\theta}] = [1, 2]$: The dashed (resp., dotted) line assumes that brand loyalty is decreasing (resp., increasing) in preferences for quality, whereas the full line, which is the efficient schedule, assumes it is constant ($\Delta t = 0$).

utility to low types (relative to the complete information outcome), leading to negative informational rents. Conversely, when $\Delta t < 0$, competition for high types is hindered by the fact that low types have strong brand loyalty, which explains why informational rents are negative for the former but positive for the latter. These conclusions generalize Proposition 1, established under binary types.

Our next result characterizes the equilibrium, and collect the findings discussed above.

Proposition 6. (Pure-strategy equilibrium: Continuum of types) *For every $\underline{t} \in (0, \underline{\tau})$, there exists $\varepsilon > 0$ such that, for all $|\Delta t| \in (0, \varepsilon)$, there exists a pure-strategy equilibrium in which the indirect utility schedule $u^*(\theta)$ satisfies, for all $\theta \in (\underline{\theta}, \bar{\theta})$, the following differential equation:*

$$\ddot{u}(\theta) = 2 - \left(\frac{1}{t(\theta)} \right) \left(\theta \dot{u}(\theta) - \frac{(\dot{u}(\theta))^2}{2} - u(\theta) \right) \quad \text{subject to} \quad \dot{u}(\theta) = \theta \quad \text{for} \quad \theta \in \{\underline{\theta}, \bar{\theta}\}. \quad (10)$$

- (a) *If $\Delta t > 0$, every interior quality involves downward distortions: $q^*(\theta) < \theta$. Moreover, relative to complete information, there are types $\theta_1, \theta_2 \in (\underline{\theta}, \bar{\theta})$ such that every type $\theta < \theta_1$ is worse-off, whereas every type $\theta > \theta_2$ is better-off.*
- (b) *If $\Delta t < 0$, every interior quality involves upward distortions: $q^*(\theta) > \theta$. Moreover, relative to complete information, there are types $\theta_1, \theta_2 \in (\underline{\theta}, \bar{\theta})$ such that every type $\theta < \theta_1$ is better-off, whereas every type $\theta > \theta_2$ is worse-off.*

The next corollary employs the characterization of Proposition 6 to perform comparative statics on both the magnitude of brand loyalty, as well as on the correlation between brand loyalty and brand preferences. To do so, it is convenient to write the brand loyalty schedule as $t(\theta) = \alpha + \beta\theta$.

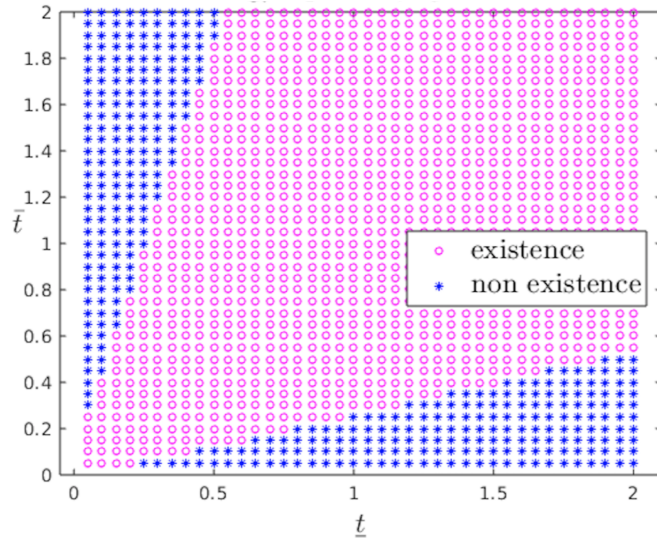


Figure 5: Brand-loyalty profiles (\underline{t}, \bar{t}) under which a pure-strategy equilibrium does (not) exist.

In this parametrization, changes in the parameter α correspond to uniform shifts on brand loyalty (across consumer types), while β captures the correlation between θ and $t(\theta)$.³⁰

Corollary 2. (Comparative Statics) *Consider the pure-strategy equilibrium of Proposition 6, and adopt the parametrization $t(\theta) = \alpha + \beta\theta$. Then, an increase in α reduces the indirect utility of every type, while an increase in β reduces the quality provision for every interior type.*

The effect of increasing α on equilibrium indirect utilities is expected. More interesting, perhaps, is that, as β increases, rendering high types more brand loyal vis-à-vis low types, quality provision decreases along the product line. These findings confirm that the comparative statics in Proposition 2 are robust to a continuum-type setting.

With a continuum of types, as in the binary-type case, a symmetric pure-strategy Nash equilibrium may fail to exist, as there is no guarantee that the firms' best responses are globally quasi-concave. We investigate this issue by numerically computing the putative pure-strategy equilibrium, and then searching for profitable incentive-compatible deviations. Assuming $[\underline{\theta}, \bar{\theta}] = [1, 2]$, and parametrizing the brand loyalty schedule by the profile (\underline{t}, \bar{t}) , Figure 5 identifies the regions where a pure-strategy equilibrium does (not) exist. The results are remarkably parallel to those under binary types: There is always a pure-strategy equilibrium close to the “diagonal” (where $\bar{t} = \underline{t}$), similarly to Proposition 1. However, non-existence obtains when brand loyalty is sufficiently different across “low” and “high” types (i.e., for \underline{t} small and \bar{t} large, or vice-versa), similarly to Remark 3. This reveals that dispersion of offers is a robust feature of competitive models involving self-selection and heterogeneous brand loyalty.

³⁰There is a one-to-one relationship between the (α, β) parametrization and the one based on the brand-loyalty profile (\underline{t}, \bar{t}) - see the proof of Proposition 6 for details. Naturally, $\beta > 0$ if and only if $\Delta t > 0$.

6.2 Timing

The design of product lines is often considered to be a long-term decision, in contrast to pricing choices, which can be altered in a short-term basis. In the baseline model, we abstracted from this distinction by assuming that prices and qualities are chosen simultaneously. It is straightforward to introduce asynchronous choices of quality and price by adopting the timing assumption first proposed by Champsaur and Roche (1989). In their model, firms are able to commit to a range of qualities before choosing prices.

Our analysis holds unchanged under this alternative timing assumption. To see why, consider first the case where the equilibrium (in the simultaneous-choice game) is in pure strategies. By committing to the set of qualities offered in the equilibrium menu, firms may implement the simultaneous-choice outcome in the asynchronous-choice game. The reason is that, ex-post, the same prices are optimal for each quality level. The same reasoning applies to mixed-strategy equilibria of the simultaneous-choice game. Firms can commit to the support of qualities offered across all equilibrium menus, to then randomize over prices (assigning a very high price to all qualities not present in the realized menu). This procedure effectively transposes every equilibrium of the simultaneous-choice game into an equilibrium of the asynchronous-choice game.

The asynchronous-choice game may however exhibit additional equilibria, as the ability to commit to a quality range generates market power (see Champsaur and Roche 1989). However, there are reasons to believe these additional equilibria are not sensible, as their existence relies on distributional assumptions (for instance, θ being uniform) and they are often discarded by refinements (of the trembling-hand type).

6.3 Consumer Participation

As discussed in Section 2, we model consumer (non-)participation by assuming that a no-purchase contract is part of the menu offered by each firm. This guarantees that the relevant margin for pricing is always the competitive one (substitution towards the competing firm), never the participation one (substitution towards the outside option). While convenient for interpretation and tractability of the model, this assumption is unessential for the conclusions of this paper.

To understand this point, consider the Hotelling model. Under the standard approach to participation (where the utility (1) is compared to zero, rather than evaluated at the no-purchase contract), the equilibrium characterization of Proposition 1 remains valid over a large set of profiles (t_l, t_h) .³¹ In region D_+ , for instance, this characterization may fail only when t_l is large (in which case u_l^* is small), as those low-type consumers located at the midpoint of the Hotelling segment would prefer not participating. Up to this point, our results on comparative statics, (non)-existence of pure-strategy

³¹The range of (t_l, t_h) for which Proposition 1 remains valid expands as θ_h gets closer to θ_l .

equilibria, and mixed-strategy equilibria remain true verbatim. Once some low types withdraw from the market under the standard modeling of participation, the parameters t_l and t_h play different roles (the former parametrizing participation incentives, the latter competition), which renders the interpretation of the model unclear. Accordingly, our modeling of participation is a convenient device to handle large values of t_l and t_h while guaranteeing full market coverage. It essentially enlarges the set of parameters under which the discrete choice framework can be fruitfully used to study competition with asymmetric information.

7 Related Literature

This article primarily contributes to the literature that studies imperfect competition in nonlinear pricing schedules (see Stole 2007 for a comprehensive survey).

In one strand of this literature, Stole (1991), Ivaldi and Martimort (1994) and Martimort and Stole (2009) study duopolistic competition in nonlinear price schedules when consumers can purchase from more than one firm. Calzolari and Denicolò (2013, 2015) evaluate the welfare impact of allowing firms to offer exclusive deals (whereby consumers get discounts if they buy nothing from the competitor). These papers speak to markets where goods are divisible, whereas our analysis is relevant for markets where purchases are inherently exclusive (e.g., most markets for durable goods).

As mentioned in the introduction, our work is more closely related to Roche and Stole (1997, 2002) and Armstrong and Vickers (2001).³² Several differences between our paper and these classic contributions stand out: First, these papers focus on the case where the propensity to switch brands is independent of one’s tastes for quality. By contrast, our focus is on the more challenging case where preferences co-move. Second, these papers adopt a “standard” Hotelling framework, in which consuming the outside option does not require incurring the transportation cost. As explained in Remark 1, this specification conflates changes in the degree of competition across firms with changes in the attractiveness of the outside option. Bypassing this limitation renders our model more tractable, while permitting a clear interpretation of comparative statics.³³ Our discrete-choice formulation is also more general, including the Hotelling setup as a special case (among others, such as Logit and Probit).

Other closely related papers are Katz (1984), Desai (2001) and Ellison (2005), who consider

³²Also relevant is the work of Stole (1995), who proposes a general model where the vertical type of each consumer interacts with her firm-specific brand taste to determine her valuation for quality. Stole compares the case where firms compete observing consumers’ brand tastes, but not vertical types, with the case where firms compete observing consumers’ vertical types, but not brand tastes. By contrast, Roche and Stole (1997, 2002) and Armstrong and Vickers (2001) assume that brand tastes do not affect one’s valuation for quality (solely determined by the vertical type) and that neither is observed by firms. We adopt this specification as well.

³³Other models exhibiting the demand specification of Roche and Stole (2002) employ numerical solution methods. Examples include Borenstein (1985), Borenstein and Rose (1994), Wilson (1993), and Yang and Ye (2008).

competitive price discrimination models related to ours, while assuming that consumers with higher valuations for quality have stronger brand preferences. Katz shows that, if downward adjacent incentive constraints bind in equilibrium, qualities have to be (weakly) under-provided. Katz however does not characterize equilibria or shows under what conditions these constraints bind. Desai provides necessary and sufficient conditions for the complete-information outcome not to violate incentive constraints.³⁴ In contrast to these contributions, we develop a full equilibrium analysis. In turn, Ellison assumes that qualities are exogenous, focusing on the equilibrium choice of prices. The emphasis of his work is in comparing two settings; one of complete information about prices, the other where consumers fail to observe the price of “upgrading” the product before getting to the store.³⁵ By contrast, the joint determination of price and quality is at the heart of the present paper, which also considers the opposite preference correlation pattern, investigates the possibility of price/quality dispersion, among many other aspects not present in Ellison’s contribution.

In turn, Bonatti (2011) develops a model of nonlinear pricing with competition where consumers’ tastes for quality are brand-specific. In this setting, conditional on choosing a given brand, high-type consumers are more brand loyal than low types, thus requiring larger discounts to switch brands. In our model, by contrast, consumers’ tastes for quality do not vary across brands, which allows us to *exogenously* change the brand loyalty of each type of consumer. Bonatti finds that quality levels are distorted downwards in equilibrium, which also occurs in our model when high-type consumers are less prone to switch brands.

More recently, Chade and Swinkels (2021) propose a model where vertically differentiated firms compete to screen consumers with private information about their willingness to pay for quality. Firms are differentiated in their ability to produce different quality levels, which leads to segmentation in equilibrium. As in here, they document the possibility of upward and downward distortions, and of non-existence of pure-strategy equilibria. We see both models as offering complementary contributions: Whereas firms are asymmetric in their ability to serve different consumer types in Chade and Swinkels (2021), the asymmetry in the current paper rather pertains to the brand loyalty of different consumer types. Accordingly, the economic rationale behind Chade and Swinkels’ findings relate to the fact that different firms sort into serving different intervals of consumer types, whereas our results hinge on the fact that firms compete for the same consumers, but do so with varying intensity.

Finally, Dessein (2003, 2004) studies competition between telecommunication networks for users with heterogenous calling patterns who self-select into their preferred calling plans. In his 2003 contribution, Dessein provides general conditions under which access charges do not affect profits.

³⁴In the language of our model, Desai’s conditions guarantee that the equilibrium lies in region E . If these conditions are violated, based on Katz’ result, Desai concludes that qualities are under-provided in equilibrium. This claim implicitly assumes that a pure-strategy equilibrium exists, which our analysis shows to be false in general.

³⁵See also Verboven (1999).

Dessein (2004) revisits this question in a setting where heavy and light users perceive the substitutability of the competing networks differently. He then shows that access charges (more often) affect profits, and provides sufficient conditions for the complete-information outcome to violate incentive constraints. Complementing Dessein’s insight, we develop a complete characterization of equilibrium.

As price discrimination is practiced under competition, consumers’ reservation values are type-dependent, in that they coincide with the indirect utility offered by the competing firms. When low types are more brand-loyal than high types (as in region D_+), the consumer’s gain in indirect utility relative to her equilibrium reservation value is decreasing in type. This is parallel to the countervailing incentives model of Lewis and Sappington (1989), and similarly leads to over-provision of quality. More generally, the analysis of firms’ best responses is related to Champsaur and Roche (1989) and Jullien (2000), who study price discrimination with type-dependent reservation utilities. However, in our model, reservation utilities are not only type-dependent but also *endogenous* and *random* (due to taste shocks). The latter renders the techniques of these papers inapplicable to our setting. The endogeneity of reservation utilities, reflected in the fixed-point nature of our problem, is also crucial for characterizing equilibrium, verifying existence in pure strategies, evaluating distortions and performing comparative statics.

Our counter-intuitive comparative statics also relate to classic contributions in price theory, where discrimination (in the form of menus) is absent. For instance, Dorfman and Steiner (1954) show that, as firms’ market power goes down (as measured by an increase in the elasticity of substitution), quality provision can either increase or decrease, as so does welfare.³⁶ More recently, Chen and Riordan (2008) show, in the context of a random utility model, that duopoly may lead to higher prices than monopoly, as increasing product variety may reduce the price-elasticity of demand. By contrast, our results are driven by the interplay between asymmetric information (manifested in the incentive constraints that shape the firms’ decisions) and competition, as captured by the varying degrees of consumers’ brand loyalty. We expect this logic to extend to more general demand systems, such as those studied by Nocke and Schutz (2018).

8 Conclusions

We study oligopolistic competition by firms practicing second-degree price discrimination. Crucially, we allow consumers with different tastes for quality to exhibit varying propensities to switch brands, which reflects a key feature of empirical models estimating demand for differentiated goods. Our analysis delivers five main take-aways:

- (a) We show how patterns of quality provision relate to co-movements between consumers’ tastes

³⁶See Dranove and Satterthwaite (2000) for a modern treatment of this seminal contribution.

for quality and brand loyalty. Specifically, quality provision is inefficiently low at the bottom (high at the top) of the product line if the propensity of low-type consumers to switch brands is small (large) relative to that of high types.

- (b) Informational rents may well be negative under competition. In fact, they are always negative for some type, while positive for the other, provided consumers obtain more than their reservation utility in equilibrium and incentive constraints bind. This is unlike the monopoly case, or the oligopoly case under uniform brand loyalty, where informational rents are always (weakly) positive.
- (c) Competition and welfare are often misaligned: More competition (in the sense that consumers become less brand-loyal) is welfare-decreasing whenever it tightens incentive constraints. This is the case, for instance, when quality provision is inefficiently high at the top of the product line. The merger from oligopoly to monopoly may then increase welfare (although it decreases consumer surplus).
- (d) Variations in the level of prices are a misleading indicator of the degree of competition in the market. In particular, more competition can either decrease or increase prices, depending on the race between quality and payoff changes.
- (e) Dispersion of offers (due to the non-existence of pure-strategy equilibria) often occurs when different types exhibit (sufficiently) different propensities to switch brands. In this case, firms randomize over ordered menus, where indirect utilities co-move across types.

Our analysis can be extended in many fruitful directions. One pertains to dynamic models of competition where consumers (with recurring consumption needs) are heterogeneous on their switching costs, while exhibiting different tastes for product characteristics (which is the case in insurance markets, for instance). Another pertains to models of competition under non-exclusive agency (as in Calzolari and Denicolò 2013), where agents differ in their tastes for quality/quantity, but also on how they perceive the complementarity/substitutability between different sellers. We expect the ideas and techniques of our paper to be useful in exploring these interesting research avenues.

9 Appendix

Proof of Proposition 0. We will verify that the strategies described in the Proposition constitute an equilibrium. Uniqueness follows from Proposition 7, which is stated and proved in the Online Appendix. We prove each claim in turn.

Claim (a). Under this putative strategy profile, no IC constraint binds and hence firm 1 solves

$$\max_{u_l^1, u_h^1} \left\{ p_l H \left(\frac{u_l^1 - u_l^*}{t} \right) (S_l^e - u_l^1) + p_h H \left(\frac{u_h^1 - u_h^*}{t} \right) (S_h^e - u_h^1) \right\}$$

subject to $u_l^1, u_h^1 \geq 0$, where (u_l^*, u_h^*) is played by every other firm. The first-order condition with respect to u_k^1 is

$$h \left(\frac{u_k^1 - u_k^*}{2t} \right) \left(\frac{S_k^e - u_k^1}{t} \right) - H \left(\frac{u_k^1 - u_k^*}{t} \right) = 0. \quad (11)$$

Notice that the problem is strictly quasi-concave by Assumption 1. After imposing symmetry, the solution (u_l^*, u_h^*) solves

$$u_k^* = S_k^e - \lambda t \geq 0,$$

as we wanted to show.

Claim (b). Under this putative strategy profile, the constraint IC_l does not bind. Hence we can write the problem as

$$\max_{u_l^1, u_h^1} \left\{ p_l H \left(\frac{u_l^1 - u_l^*}{t} \right) (S_l(u_l^1, u_h^1) - u_l^1) + p_h H \left(\frac{u_h^1 - u_h^*}{t} \right) (S_h^e - u_h^1) \right\}.$$

In this case, the first-order condition with respect to u_l^1 is

$$\begin{aligned} p_l h \left(\frac{u_l^1 - u_l^*}{t} \right) \left(\frac{S_l(u_l^1, u_h^1) - u_l^1}{t} \right) + p_l H \left(\frac{u_l^1 - u_l^*}{t} \right) \left(\frac{\partial S_l(u_l^1, u_h^1)}{\partial u_l} - 1 \right) \\ \leq p_l h \left(\frac{u_l^1}{t} \right) \left(\frac{S_l^e - u_l^1}{t} \right) - p_l H \left(\frac{u_l^1}{t} \right). \end{aligned}$$

Since $t > \frac{\bar{\eta}}{\lambda}$, the expression above is nonpositive at $u_l^1 = 0$, which is then optimal by strict quasi-concavity. We conclude that $u_l^* = 0$. Imposing this, the first-order condition with respect to u_h^1 is then

$$p_l H(0) \left(\frac{\partial S_l(0, u_h^1)}{\partial u_h^1} \right) + p_h h \left(\frac{u_h^1 - u_h^*}{2t} \right) \left(\frac{S_h^e - u_h^1}{t} \right) - p_h H \left(\frac{u_h^1 - u_h^*}{t} \right).$$

First assume that

$$p_l H(0) \left(\frac{\partial S_l(0, 0)}{\partial u_h^1} \right) + p_h h(0) \left(\frac{S_h^e}{t} \right) - p_h H(0) \leq 0,$$

and observe that

$$p_h h(0) \left(\frac{S_h^e}{t} \right) - p_h H(0) \leq -p_l H(0) \left(\frac{\partial S_l(0, 0)}{\partial u_h^1} \right) < 0, \quad (12)$$

which will be useful below. We will show that if all other firms choose $u_h^* = 0$ then choosing $u_h^1 = 0$ is optimal. Rearrange the first-order condition at any putative best response u_h^1 as

$$p_h H(u_h^1) \left(\frac{h(u_h^1)}{H(u_h^1)} \left(\frac{S_h^e - u_h^1}{t} \right) - 1 \right) + p_l H(0) \left(\frac{\partial S_l(0, u_h^1)}{\partial u_h} \right) \leq 0.$$

Differentiating it we obtain

$$\begin{aligned} p_h h(u_h^1) \left(\frac{h(u_h^1)}{H(u_h^1)} \left(\frac{S_h^e - u_h^1}{t} \right) - 1 \right) + p_h h(u_h^1) \left(\frac{\partial}{\partial u_h} \frac{h(u_h^1)}{H(u_h^1)} \right) \left(\frac{S_h^e - u_h^1}{t} \right) \\ - p_h h(u_h^1) \left(\frac{S_h^e}{t} \right) + p_l H(0) \left(\frac{\partial S_l^2(0, u_h^1)}{\partial u_h^2} \right) < 0, \end{aligned}$$

where the first term is negative by (12), the second because H is strictly log-concave and the last because $\left(\frac{\partial S_l^2(0, u_h^1)}{\partial u_h^2} \right) < 0$. This shows that whenever the first-order condition holds we have a local maximum. This implies that there cannot be any other critical point and guarantees that $u_h^* = 0$ is optimal and hence $(u_l^*, u_h^*) = (0, 0)$ is an equilibrium.

Next assume that

$$p_h H(0) \left(\frac{h(0)}{H(0)} \left(\frac{S_h^e}{t} \right) - 1 \right) + p_l H(0) \left(\frac{\partial S_l(0, 0)}{\partial u_h} \right) > 0,$$

and define $u_h^* > 0$ as the solution to

$$\frac{S_h^e - u_h^*}{\lambda t} + \frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, u_h^*) - 1 = 0. \quad (13)$$

By a similar argument to the one above we conclude that $(0, u_h^*)$ is an equilibrium. Next assume that there is $t > 0$ such that $u_h^* = 0$. Then $\frac{S_h^e}{\lambda t} + \frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, 0) - 1 \leq 0$ and $\frac{S_h^e}{\lambda t} + \frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, 0) - 1 \leq 0$ for all $\tilde{t} > t$. Moreover, $\frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, 0) - 1 < 0$, which shows that the monopolistic solution displays $(0, 0)$. Next assume that $u_h^* > 0$ for every t . Using (13) we have

$$\frac{\partial u_h^*}{\partial t} = \frac{\frac{S_h^e - u_h^*}{\lambda t^2}}{\frac{p_l}{p_h} \frac{\partial^2 S_l}{\partial u_h^2}(0, u_h^*) - 1} < 0$$

and by a straightforward argument we see that u_h^* converges to the monopolistic level u_h^∞ , which solves

$$\frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(0, u_h^\infty) - 1 = 0.$$

This concludes the proof. ■

Proof of Proposition 1. We prove each claim in turn.

Claim (a). Assume that $(t_l, t_h) \in E$. Consider the problem in which both IC constraints are ignored. Firm 1 solves

$$\max_{u_l^1, u_h^1} \left\{ p_l H \left(\frac{u_l^1 - u_l^*}{t_l} \right) (S_l^e - u_l^1) + p_h H \left(\frac{u_h^1 - u_h^*}{t_h} \right) (S_h^e - u_h^1) \right\}$$

subject to $u_l^1, u_h^1 \geq 0$, where (u_l^*, u_h^*) is played by every other firm. The problem is separable and strictly quasi-concave with relation to the utility provided to each one of the types. Imposing symmetry, the first-order conditions with respect to u_k , $k \in \{l, h\}$, are

$$S_k^e - u_k^* - \lambda t_k \leq 0, \quad (14)$$

with equality if $u_k^* > 0$. The unique solution to (14) is given by $u_k^* = \max\{S_k^e - \lambda t_k, 0\}$, $k \in \{l, h\}$. Since the problem is separable, this strategy profile is optimal whenever neither constraints binds, which holds because $(t_l, t_h) \in E$. Finally uniqueness follows from Proposition 7, which is stated and proved in the Online Appendix.

Claim (b). Assume that $(t_l, t_h) \in D_+$. Proposition 7 (Online Appendix) establishes that there exists at most one pure-strategy equilibrium. Consider a pure-strategy equilibrium. Assume towards a contradiction that there exists a pure-strategy equilibrium (u_h^*, u_l^*) in which no IC constraint binds, in which case $u_h^* - u_l^* \in [\Delta\theta q_l^e, \Delta\theta q_h^e]$. In this case, we have $u_k^* = \max\{S_k^e - \lambda t_k, 0\}$, $k \in \{l, h\}$ and hence

$$\begin{aligned} \Lambda(t_l, t_h) &\equiv \lambda t_h + \max\{S_l^e - \lambda t_l, 0\} \\ &\geq S_h^e - u_h^* + u_l^* \\ &\geq S_h^e - q_h^e \Delta\theta = \eta_h, \end{aligned}$$

a contradiction. Next assume that there exists a pure-strategy equilibrium in which IC_h binds and hence $u_h^* - u_l^* < \Delta\theta q_l^e$. The first-order condition with respect to u_h and u_l read

$$\begin{aligned} S_h^e - u_h^* + \lambda t_h \left(\frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h}(u_l^*, u_h^*) - 1 \right) &\leq 0, \\ S_l(u_l^*, u_h^*) - u_l^* + \lambda t_l \left(\frac{\partial S_l}{\partial u_l}(u_l^*, u_h^*) - 1 \right) &\leq 0. \end{aligned}$$

with equality $u_h^* > 0$ and $u_l^* > 0$ respectively. From the first-order condition we have $\lambda t_h \geq S_h^e - u_h^*$. If $u_l^* = 0$, then $\max\{S_l^e - \lambda t_l, 0\} \geq u_l^*$. Otherwise,

$$S_l^e - \lambda t_l > S_l(u_l^*, u_h^*) + \lambda t_l \left(\frac{\partial S_l}{\partial u_l}(u_l^*, u_h^*) - 1 \right) = u_l^*.$$

Therefore we have

$$\Lambda(t_l, t_h) \geq S_h^e - u_h^* + u_l^* > S_h^e - q_l^e \Delta\theta,$$

a contradiction.

Next, define $\tau_h^e(t_l)$ by

$$\lambda \tau_h^e(t_l) + \max\{S_l^e - \lambda t_l, 0\} = \eta_h.$$

Hence if $t_h < \tau_h^e(t_l)$ we have $\Lambda(t_l, t_h) < \eta_h$ and hence $(t_l, t_h) \in D_+$. We will show that there exists $\xi > 0$ such if $t_h \in (\tau_h^e(t_l) - \xi, \tau_h^e(t_l))$ then there is a symmetric pure-strategy equilibrium in which (u_l^*, u_h^*) solve

$$S_h(u_l^*, u_h^*) - u_h^* + \lambda t_h \left(\frac{\partial S_h}{\partial u_h}(u_l^*, u_h^*) - 1 \right) = 0 \quad (15)$$

and

$$\frac{S_l^e - u_l^*}{\lambda t_l} + \left(\frac{p_h}{p_l} \frac{\partial S_h}{\partial u_l}(u_l^*, u_h^*) - 1 \right) \leq 0. \quad (16)$$

First start at $(t_l, \tau_h^e(t_l))$ and notice that the pure strategy equilibrium satisfies both (15) and (16) with $S_h(u_l^*, u_h^*) = S_h^e$ and $\frac{\partial S_h}{\partial u_l}(u_l^*, u_h^*) = 0$.

Assume first that $u_l^* > 0$. Write $\Pi((u_l, u_h) | (u_l^*, u_h^*), (t_l, t_h))$ for the profit of a firm when other firms choose (u_l^*, u_h^*) and parameters are (t_l, t_h) . Assumption 1 implies that

$$\frac{\partial^2 \Pi((u_l^*, u_h^*) | (u_l^*, u_h^*), (t_l, t_h))}{\partial^2 u_k} < 0$$

for $k = l, h$. Moreover, $S_h(u_l^*, u_h^*) = S_h^e$ implies that

$$\frac{\partial^2 \Pi((u_l^*, u_h^*) | (u_l^*, u_h^*), (t_l, t_h))}{\partial u_l \partial u_h} = 0.$$

We conclude that the matrix $\partial^2 \Pi((u_l^*, u_h^*) | (u_l^*, u_h^*), (t_l, t_h))$ is locally negative definite. Moreover, Assumption 1 also implies that (u_l^*, u_h^*) is the unique maximizer of

$$(u_l, u_h) \rightarrow \Pi((u_l, u_h) | (u_l^*, u_h^*), (t_l, t_h)).$$

By the implicit function theorem and a by a straightforward continuity argument, there exists $\xi > 0$ such that, if $\tilde{t}_h \in (\tau_h^e(t_l) - \xi, \tau_h^e(t_l))$, then the solution to

$$S_h(u_l^*(\tilde{t}_h), u_h^*(\tilde{t}_h)) - u_h^*(\tilde{t}_h) + \lambda t_h \left(\frac{\partial S_h}{\partial u_h}(u_l^*(\tilde{t}_h), u_h^*(\tilde{t}_h)) - 1 \right) = 0$$

and

$$\frac{S_l^e - u_l^*(\tilde{t}_h)}{\lambda t_l} + \left(\frac{p_h}{p_l} \frac{\partial S_h}{\partial u_l}(u_l^*(\tilde{t}_h), u_h^*(\tilde{t}_h)) - 1 \right) = 0$$

is a pure-strategy equilibrium. Moreover, notice that $(t_l, \tilde{t}_h) \in D_+$ and recall that we have ruled out above the possibility that no IC binds and that IC_h binds. Hence we conclude that IC_l binds.

Finally, if $u_l^* = 0$, one can use a similar argument to show that there exists a $\xi > 0$ such that $\tilde{t}_h \in (\tau_h^e(t_l) - \xi, \tau_h^e(t_l))$ then the equilibrium will be given by

$$S_h(0, u_h^*(\tilde{t}_h)) - u_h^*(\tilde{t}_h) + \lambda t_h \left(\frac{\partial S_h}{\partial u_h}(0, u_h^*(\tilde{t}_h)) - 1 \right) = 0 \quad (17)$$

and

$$\frac{S_l^e}{\lambda t_l} + \left(\frac{p_h}{p_l} \frac{\partial S_h}{\partial u_l}(0, u_h^*(\tilde{t}_h)) - 1 \right) \leq 0, \quad (18)$$

with $u_h^*(\tilde{t}_h)$ being a solution to (17). By the same reason as above, we conclude that IC_l binds.

Claim (c). Assume that $(t_l, t_h) \in D_-$. The proof is analogous to the one of claim (b) and therefore omitted. ■

Proof of Proposition 2. We proceed by analyzing five different cases.

Case 1 Assume that (t_l, t_h) belong to an open set $U \in E$.

In this case, we have $q_k = q_k^e$ and since $u_k = \max \{S_k^e - \lambda t_k, 0\}$, $\frac{\partial u_k}{\partial t_h} = -\lambda \{S_k^e - \lambda t_k, 0\}$.

Case 2 Assume that (t_l, t_h) belong to an open set $U \in D_-$.

Case 2.1 Assume that IC_h binds and IR does not in U .

The equilibrium is given by the following equations:

$$\begin{aligned} S_l(u_l, u_h) + \lambda t_l \left(\frac{\partial S_l(u_l, u_h)}{\partial u_l} - 1 \right) &= u_l \\ S_h^e - \lambda t_h + \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(u_l, u_h)}{\partial u_h} &= u_h. \end{aligned}$$

Let $\Delta(u_l, u_h) := u_h - u_l$. To ease notation, we write only Δ below, leaving implicit its dependence on (u_l, u_h) wherever it does not lead to confusion. Consider the equation

$$G(\Delta, t_l, t_h) \equiv \left[S_h^e - \lambda t_h + \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(\Delta)}{\partial u_h} \right] - \left[S_l(\Delta) + \lambda t_l \left(\frac{\partial S_l(\Delta)}{\partial u_l} - 1 \right) \right] - \Delta.$$

We have

$$\begin{aligned} \frac{\partial G(\Delta, t_l, t_h)}{\partial \Delta} &= \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial^2 S_l(\Delta)}{\partial \Delta^2} - \lambda \frac{\partial S_l(\Delta)}{\partial \Delta} + \lambda t_l \frac{\partial^2 S_l(\Delta)}{\partial \Delta^2} < 0 \\ \frac{\partial G(\Delta, t_l, t_h)}{\partial t_h} &= -\lambda + \lambda \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(\Delta)}{\partial \Delta} = -\left(\frac{S_h^e - u_h}{t_h} \right) < 0 \\ \frac{\partial G(\Delta, t_l, t_h)}{\partial t_l} &= \lambda \left(1 - \frac{\partial S_l(\Delta)}{\partial u_l} \right) = \left(\frac{S_l(u_l, u_h) - u_l}{t_l} \right) > 0 \end{aligned}$$

Therefore by straightforward algebra we have $\frac{\partial \Delta}{\partial t_l} > 0$, $\frac{\partial \Delta}{\partial t_h} < 0$, $\frac{\partial u_h}{\partial t_l} = \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial^2 S_l(u_l, u_h)}{\partial u_h^2} \frac{\partial \Delta}{\partial t_l} < 0$, $\frac{\partial u_h}{\partial t_h} = \frac{\partial \Delta}{\partial t_h} + \frac{\partial u_l}{\partial t_h} < 0$, $\frac{\partial u_l}{\partial t_h} = \left[\frac{\partial S_l(u_l, u_h)}{\partial \Delta} - \lambda t_l \left(\frac{\partial^2 S_l(u_l, u_h)}{\partial \Delta^2} \right) \right] \frac{\partial \Delta}{\partial t_h} < 0$ and $\frac{\partial u_l}{\partial t_l} = \frac{\partial u_h}{\partial t_l} - \frac{\partial \Delta}{\partial t_l} < 0$.

Case 2.2 Assume that IC_h and IR bind and $u_h > 0$ in U .

The equilibrium is given by the following equations:

$$\begin{aligned} u_l &= 0 \\ S_h^e - \lambda t_h + \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(0, u_h)}{\partial u_h} &= u_h. \end{aligned}$$

Locally, we have $\frac{\partial \Delta}{\partial t_l} = 0$, $\frac{\partial \Delta}{\partial t_h} = -\frac{\lambda \left(1 - \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(0, u_h)}{\partial u_h} \right)}{\left(1 - \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial^2 S_l(0, u_h)}{\partial u_h^2} \right)} < 0$, $\frac{\partial u_h}{\partial t_l} = 0$, $\frac{\partial u_h}{\partial t_h} = \frac{\partial \Delta}{\partial t_h} < 0$, $\frac{\partial u_l}{\partial t_h} = 0$ and

$$\frac{\partial u_l}{\partial t_l} = 0.$$

Case 2.3 Assume that $u_h = 0$ in U .

In this case, $q_l = 0 = u_l = y_l = 0$, $q_h = q_h^e$.

Case 3 Assume that (t_l, t_h) belong to an open set $U \in D_+$.

Case 3.1 Assume that IC_l binds and IR does not in U .

The equilibrium is given by the following equations:

$$\begin{aligned} u_l &= S_l^e + \lambda t_l \left(\frac{p_h}{p_l} \right) \left(\frac{\partial S_h(u_l, u_h)}{\partial u_l} \right) - \lambda t_l \\ u_h &= S_h(u_l, u_h) + \lambda t_h \frac{\partial S_h(u_l, u_h)}{\partial u_h} - \lambda t_h \end{aligned}$$

Consider the equation

$$G(\Delta, t_l, t_h) := \left[S_h(u_l, u_h) - \lambda t_h + \lambda t_h \frac{\partial S_h(\Delta)}{\partial u_h} \right] - \left[S_l^e + \lambda t_l \left(\left(\frac{p_h}{p_l} \right) \frac{\partial S_h(\Delta)}{\partial u_l} - 1 \right) \right] - \Delta.$$

We have:

$$\begin{aligned} \frac{\partial G(\Delta, t_l, t_h)}{\partial \Delta} &= \left[\frac{\partial S_h(u_l, u_h)}{\partial \Delta} + \lambda t_h \frac{\partial^2 S_h(\Delta)}{\partial \Delta^2} + \lambda t_l \left(\frac{p_h}{p_l} \right) \frac{\partial^2 S_h(\Delta)}{\partial \Delta^2} \right] - 1 < 0 \\ \frac{\partial G(\Delta, t_l, t_h)}{\partial t_l} &= -\lambda \left(\left(\frac{p_h}{p_l} \right) \frac{\partial S_h(\Delta)}{\partial u_l} - 1 \right) = \frac{1}{t_l} (S_l^e - u_l) > 0 \\ \frac{\partial G(\Delta, t_l, t_h)}{\partial t_h} &= -\lambda \left(1 - \frac{\partial S_h(\Delta)}{\partial u_h} \right) = -\frac{1}{t_h} (S_h(u_l, u_h) - u_h) < 0 \end{aligned}$$

Therefore we have $\frac{\partial \Delta}{\partial t_l} > 0$, $\frac{\partial \Delta}{\partial t_h} < 0$, $\frac{\partial u_h}{\partial t_l} = \left(\lambda t_h \frac{\partial^2 S_h(\Delta)}{\partial u_h^2} + \frac{\partial S_h(\Delta)}{\partial u_h} \right) \frac{\partial \Delta}{\partial t_l} < 0$, $\frac{\partial u_h}{\partial t_h} = \frac{\partial \Delta}{\partial t_h} + \frac{\partial u_l}{\partial t_h} < 0$, $\frac{\partial u_l}{\partial t_h} = -\lambda t_l \left(\frac{p_h}{p_l} \right) \left(\frac{\partial S_h^2(u_l, u_h)}{\partial \Delta^2} \right) \frac{\partial \Delta}{\partial t_h} < 0$ and $\frac{\partial u_l}{\partial t_l} = \frac{\partial u_h}{\partial t_l} - \frac{\partial \Delta}{\partial t_l} < 0$.

Case 3.2 Assume that IC_l and IR binds in U .

The equilibrium is given by the following equations:

$$\begin{aligned} u_l &= 0 \\ u_h &= S_h(0, u_h) - \lambda t_h + \lambda t_h \frac{\partial S_h(0, u_h)}{\partial u_h} \end{aligned}$$

We have $\frac{\partial \Delta}{\partial t_l} = 0$, $\frac{\partial \Delta}{\partial t_h} = -\frac{\lambda \left(1 - \frac{\partial S_h(0, u_h)}{\partial u_h} \right)}{\left(1 - \frac{\partial S_h(0, u_h)}{\partial u_h} - \lambda t_h \frac{\partial^2 S_h(0, u_h)}{\partial u_h^2} \right)} < 0$, $\frac{\partial u_h}{\partial t_l} = 0$, $\frac{\partial u_h}{\partial t_h} = \frac{\partial \Delta}{\partial t_h} < 0$, $\frac{\partial u_l}{\partial t_l} = 0$ and $\frac{\partial u_l}{\partial t_h} = 0$.

This exhausts all cases and completes the proof. ■

References

- [1] Anderson, S., A. de Palma, and J.-F. Thisse, 1992. *Discrete Choice Theory of Product Differentiation*, MIT Press.
- [2] Armstrong, M. and J. Vickers, 2001. "Competitive Price Discrimination." *Rand Journal of Economics*, 32, 579-605.
- [3] Baye, M., and J. Morgan, 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *American Economic Review*, 91(3): 454-474.
- [4] Berry, S., J. Levinsohn and A. Pakes, 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63, 841-890.
- [5] Bénabou, R. and J. Tirole, 2016. "Bonus Culture: Competitive Pay, Screening, and Multitasking." *Journal of Political Economy*, 124(2), 305-370.
- [6] Bonatti, A., 2011. "Brand-Specific Tastes for Quality." *International Journal of Industrial Organization*. 29, 562-575.
- [7] Borenstein, S., 1985. "Price Discrimination in Free-Entry Markets." *Rand Journal of Economics*, 16, 380-397.
- [8] Borenstein, S., and N. Rose, 1994. "Competition and Price Dispersion in the U.S. Airline Industry." *Journal of Political Economy*, 102, 653-683.

- [9] Burdett, K., and D. Judd, 1983. "Equilibrium Price Dispersion." *Econometrica*, 51, 955-970.
- [10] Calzolari, G., and V. Denicolò, 2013. "Competition with Exclusive Contracts and Market-Share Discounts." *American Economic Review*, 103, 2384-2411.
- [11] Calzolari, G., and V. Denicolò, 2015. "Exclusive Contracts and Market Dominance." *American Economic Review*, 105, 3321-51.
- [12] Chade, H., and J. Swinkels, 2021. "Screening in Vertical Oligopolies." *Econometrica*, 89(3), 1265-1311.
- [13] Champsaur, P., and J.-C. Roche, 1989. "Multi-Product Duopolists". *Econometrica*, 57, 533-557.
- [14] Chen, Y., and M. Riordan, 2007. "Price and Variety in the Spokes Model." *The Economic Journal*, 117(522), 897-921.
- [15] Chen, Y., and M. Riordan, 2008. "Price-Increasing Competition." *Rand Journal of Economics*, 39(4), 1042-1058.
- [16] Chu, C. S., 2010. "The Effect of Satellite Entry on Cable Television Prices and Product Quality." *Rand Journal of Economics*, 41, 730-764.
- [17] Crawford, G., 2012. "Accommodating Endogenous Product Choices: A Progress Report." *International Journal of Industrial Organization*, 30, 315-320.
- [18] Crawford, G., O. Shcherbakov, and M. Shum, 2019. "Quality Overprovision in Cable Television Markets." *American Economic Review*, 109(3), 956-995.
- [19] Desai, P., 2001. "Quality Segmentation in Spatial Markets: When Does Cannibalization Affect Product Line Design?" *Marketing Science*, 20(3), 265-283.
- [20] Dessein, W., 2003. "Network Competition in Nonlinear Pricing." *Rand Journal of Economics*, 34(4), 593-611.
- [21] Dessein, W., 2004. "Network Competition with Heterogeneous Customers and Calling Patterns." *Information Economics and Policy*, 16, 323-345.
- [22] Dorfman, R., and P. Steiner, 1954. "Optimal advertising and optimal quality." *American Economic Review*, 44(5), 826-836.
- [23] Dranove, D., and M. Satterthwaite, 2000. "The industrial organization of health care markets." In Culyer, A. and Newhouse, J., editors, *Handbook of Health Economics*, chapter 20, pages 1094-1139. Elsevier Science, North-Holland, New York and Oxford.
- [24] Dupuit, J., 1849. "De l'Influence des Péages sur l'Utilité des Voies de Communication", *Annales des Ponts et Chaussées*, s. II, 1st semester, pp. 170-248; English translation of the last section "On Tolls and Transport Charges", by E. Henderson, *International Economic Papers* (1962), n. 11, pp. 7-31.
- [25] Durrmeyer, I., 2020. "Winners and Losers: The Distributional Effects of the French Feebate on the Automobile Market." Working Paper.
- [26] Ellison, G., 2005. "A Model of Add-on Pricing." *Quarterly Journal of Economics*, 120, 585-367.
- [27] Fan, Y., 2013. "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market." *American Economic Review*, 103(5), 1598-628.

- [28] Gabaix, X., 2014. “A Sparsity-Based Model of Bounded Rationality.” *Quarterly Journal of Economics*, 129(4), 1661-1710.
- [29] Garrett, D., R. Gomes and L. Maestri (2019): “Competitive Screening under Heterogeneous Information.” *Review of Economic Studies*, 86, 1590-1630.
- [30] Gandhi, A., L. Froeb, S. Tschantz and G. Werden (2008): “Post-Merger Product Repositioning.” *The Journal of Industrial Economics*, 56, 49-67.
- [31] Gomes, R., J.-M. Lozachmeur and L. Maestri (2020): “Nonlinear Pricing in Oligopoly: How Brand Preferences Shape Market Outcomes.” CEPR Discussion Paper DP15253.
- [32] Ivaldi, M., and D. Martimort, 1994. “Competition Under Nonlinear Pricing.” *Annales d’Économie et de Statistique*, 34, 71-114.
- [33] Jullien, B., 2000. “Participation Constraints in Adverse Selection Models”. *Journal of Economic Theory*, 93, 1-47.
- [34] Kaplan, G., and G. Menzio (2015). “The Morphology of Price Dispersion”, *International Economic Review*, 56(4), 1165-1204.
- [35] Kaplan, G., G. Menzio, L. Rudanko, and N. Trachter (2019). “Relative Price Dispersion: Evidence and Theory”, *American Economic Journal: Microeconomics*, 11(3), 68-124.
- [36] Katz, M., 1984. “Firm-Specific Differentiation and Competition Among Multiproduct Firms”, *The Journal of Business*, 57(1), S149-S166.
- [37] Lewis, T., and D. Sappington, 1988. “Countervailing Incentives in Agency Problems”, *Journal of Economic Theory*, 49, 294-313.
- [38] Maskin, E., and J. Riley, 1984. “Monopoly with Incomplete Information.” *Rand Journal of Economics*, 15, 171-196.
- [39] Martimort, D., and L. Stole, 2009. “Market Participation in Delegated and Intrinsic Common-Agency Games.” *Rand Journal of Economics* 40, 78-102.
- [40] McManus, B., 2007. “Nonlinear pricing in an oligopoly market: the case of specialty coffee.” *Rand Journal of Economics*, 38, 512-532.
- [41] Montez, J., and N. Schutz, 2020. “All-Pay Oligopolies under Incomplete Information.” Working paper.
- [42] Mussa, M. and S. Rosen, 1978. “Monopoly and Product Quality.” *Journal of Economic Theory*, 18, 301-317.
- [43] Nash, J., 1950. “Equilibrium Points in n-Person Games.” *Proceedings National Academy Sciences*, 36, 48-49.
- [44] Nevo, A., 2001. “Measuring Market Power in the Ready-to-Eat Cereal Industry.” *Econometrica*, 69(2), 307-342.
- [45] Nielsen, 2015. “The Total Audience Report: Q3 2015.”
- [46] Nocke, V., and N. Schutz, 2018. “Multi-Product Firm Oligopoly: An Aggregative Games Approach.” *Econometrica*, 86(2), 523-557.
- [47] Petrin, A., 2002. “Quantifying the Benefits of New Products: The Case of the Minivan.” *Journal of Political Economy*, 110(4), 705-729.

- [48] Roche, J.-C., and L. Stole, 1997. "Competitive Nonlinear Pricing." Working paper.
- [49] Roche, J.-C., and L. Stole, 2002. "Nonlinear Pricing with Random Participation." *Review of Economic Studies*, 69, 277-311.
- [50] Shelegia, S., and C. Wilson, 2020. "A Generalized Model of Advertised Sales." *American Economic Journal: Microeconomics*, forthcoming.
- [51] Stole, L., 1991. "Mechanism Design and Common Agency". Working Paper.
- [52] Stole, L., 1995. "Nonlinear Pricing and Oligopoly." *Journal of Economics & Management Strategy*, 4, 529-562.
- [53] Stole, L., 2007. "Price Discrimination and Competition." In: *Handbook of Industrial Organization*, vol. 3 (eds. M. Armstrong and R. Porter), North-Holland.
- [54] Varian, H., 1980. "A Model of Sales." *American Economic Review*, 70, 651-659.
- [55] Varian, H., 2006. *Intermediate microeconomics*. 7th ed. New York: Norton.
- [56] Verboven, F., 1999. "Product Line Rivalry and Price Discrimination." *Journal of Industrial Economics*, 47, 399-426.
- [57] Wilson, R., 1993. *Nonlinear Pricing*. Oxford University Press, Oxford.
- [58] Yang, H., and Ye, L., 2008. "Nonlinear Pricing, Market Coverage and Competition." *Theoretical Economics*, 3, 123-153.

Online Appendix

Omitted Proofs

Proof of Corollary 1. Recall from Proposition 0 that the monopolistic (Mussa-Rosen) solution satisfies

$$\frac{\partial S_l(0, u_h^\infty)}{\partial u_h} = \frac{p_h}{p_l}.$$

If $(t_l, t_h) \in D_-$, whenever $u_h > 0$, Proposition 1 implies that

$$\frac{\partial S_l(u_l, u_h)}{\partial u_h} = \frac{p_h}{p_l} - \frac{S_h^e - u_h p_h}{\lambda t_h} \frac{p_h}{p_l} < \frac{p_h}{p_l} = \frac{\partial S_l}{\partial u_h}(0, u_h^\infty),$$

which implies that $q_l(u_l, u_h) > q_h^\infty$, as the function $S_l(u_l, u_h)$ is concave in $u_h - u_l$. Hence the merger is always welfare-reducing in this region.

Now let $(t_l, t_h) \in D_+$. By Proposition 2, $\frac{\partial q_h}{\partial t_h} \leq 0$, which implies that equilibrium welfare is the lowest under oligopoly when $t_h = 0$. Proposition 4 establishes that, when $t_h = 0$, a unique pure-strategy equilibrium exists provided $t_l \geq \tilde{t}_l$, where

$$\tilde{t}_l \equiv \inf \left\{ t_l : \frac{\partial S_h}{\partial u_l}(0, \hat{u}_h) < \frac{1}{J} \frac{p_l}{p_h} \left(1 - \frac{S_l^e}{\lambda t_l} \right) \right\}.$$

This equilibrium is such that $u_l^* = 0$ and $u_h^* = \hat{u}_h$, where \hat{u}_h implicitly solves $S_h(0, \hat{u}_h) = \hat{u}_h$. Welfare then equals

$$W(\tilde{t}_l, 0) \equiv p_l S_l^e + p_h \hat{u}_h.$$

Under monopoly, welfare is

$$W^\infty \equiv p_l S_l(0, u_h^\infty) + p_h S_h^e.$$

To show that welfare under oligopoly may be lower than under monopoly, suppose the production cost is quadratic: $\varphi(q) = \frac{1}{2}q^2$. Then $\hat{u}_h = 2\theta_l \Delta\theta$, and

$$\tilde{t}_l = \frac{S_l^*}{\lambda \kappa_l} = \frac{S_l^*}{\lambda \left(1 - 2 \frac{p_h}{p_l} \left(\frac{2\theta_l - \theta_h}{\Delta\theta} \right) \right)},$$

where $\kappa_l = 1 - 2 \frac{p_h}{p_l} \left(\frac{2\theta_l - \theta_h}{\Delta\theta} \right) > 0$ guarantees that $\tilde{t}_l < \infty$.

Welfare is then

$$W(\tilde{t}_l, 0) = p_l S_l^e + p_h 2\theta_l \Delta\theta.$$

Under monopoly, simple algebra reveals that

$$S_l(0, u_h^\infty) = \frac{1}{2} \left((\theta_l)^2 - \left(\frac{p_h}{p_l} \Delta\theta \right)^2 \right),$$

in which case

$$W^\infty = p_l \left(\frac{1}{2} \left((\theta_l)^2 - \left(\frac{p_h}{p_l} \Delta\theta \right)^2 \right) \right) + p_h \frac{1}{2} (\theta_h)^2.$$

It then follows from simple computations that

$$W^\infty > W(\tilde{t}_l, 0) \iff \left(\frac{\theta_h}{\Delta\theta} \right)^2 > 4 \frac{\theta_l}{\Delta\theta} + \frac{p_h}{p_l}.$$

We should compare $W(\tilde{t}_l, 0)$ and W^∞ under the following restrictions: $2\theta_l > \theta_h$, which is Assumption 2, $\kappa_l > 0$, which assures $\tilde{t}_l < \infty$, and $\theta_l - \frac{p_h}{p_l} \Delta\theta > 0$, which assures that $S_l(0, u_h^\infty) > 0$. Notice that, by taking p_l to one, the inequalities $\kappa_l > 0$ and $\theta_l - \frac{p_h}{p_l} \Delta\theta > 0$ are trivially satisfied. Hence it suffices to find $\alpha \in (\frac{1}{2}, 1)$ such that $\theta_l = \alpha\theta_h$ and

$$\left(\frac{\theta_h}{\Delta\theta} \right)^2 > 4 \frac{\theta_l}{\Delta\theta} \iff \left(\frac{1}{1-\alpha} \right) \left(\frac{1}{1-\alpha} \right) > 4 \frac{\alpha}{1-\alpha},$$

which holds if and only if

$$\left(\frac{1}{\alpha(1-\alpha)} \right) > 4,$$

which holds for every $\alpha > \frac{1}{2}$. In other words, whenever $2\theta_l > \theta_h$, there is $\hat{p}_l \in (0, 1)$ such that $p_l > \hat{p}_l$ implies that $W^\infty > W(\tilde{t}_l, 0)$. \blacksquare

Proof of Proposition 3. We proceed by analyzing five different cases.

Case 1 Assume that (t_l, t_h) belong to an open set $U \in E$.

We have $u_k = \max\{S_k^e - t_k, 0\}$, $k \in \{l, h\}$, which immediately implies $\frac{\partial y_k}{\partial t_k} = \mathbb{I}_{\{S_k^e - t_k > 0\}}$.

Case 2 Assume that (t_l, t_h) belong to an open set $U \in D_+$.

Case 2.1 Assume that IC_l binds and IR does not in U .

The equilibrium is given by the following equations:

$$\begin{aligned} S_l^e + \lambda t_l \left(\frac{p_h}{p_l} \right) \left(\frac{\partial S_h(u_l, u_h)}{\partial u_l} \right) - \lambda t_l &= u_l \\ S_h(u_l, u_h) + \lambda t_h \frac{\partial S_h(u_l, u_h)}{\partial u_h} - \lambda t_h &= u_h. \end{aligned}$$

We have $y_h = \theta_h q_h(u_l, u_h) - u_h$ and $y_l = \theta_l q_l^e - u_l$.

We have the following comparative statics: $\frac{\partial y_l}{\partial t_l} = -\frac{\partial u_l}{\partial t_l} > 0$, $\frac{\partial y_l}{\partial t_h} = -\frac{\partial u_l}{\partial t_h} > 0$, $\frac{\partial y_h}{\partial t_l} = \left(\frac{\theta_h}{\Delta\theta} \right) \frac{\partial \Delta}{\partial t_l} - \frac{\partial u_h}{\partial t_l} > 0$ and

$$\frac{\partial y_h}{\partial t_h} = \frac{\partial \Delta}{\partial t_h} \left(\frac{1}{\Delta\theta} \right)^2 \left[\theta_l \Delta\theta - \lambda t_l \left(\frac{p_h}{p_l} \right) \varphi''(q_h(u_l, u_h)) \right].$$

The desired quasi-convexity when $\varphi''' \leq 0$ follows because $\text{sign} \left[\frac{\partial y_h}{\partial t_h} \right] = \text{sign} \left[t_l \left(\frac{p_h}{p_l} \right) \varphi''(q_h) - \theta_l \Delta\theta \right]$. Since $\frac{\partial q_h}{\partial t_h} < 0$, we conclude that $\frac{\partial y_h}{\partial t_h} = 0$ implies $\frac{\partial y_h}{\partial t_h} \geq 0$ for every $\tilde{t}_h > t_h$.

Case 2.2 Assume that IC_l and IR binds in U .

The equilibrium is given by the following equation:

$$S_h(0, u_h) + \lambda t_h \frac{\partial S_h(0, u_h)}{\partial u_h} - \lambda t_h = u_h.$$

Prices are given by $y_h = \theta_h q_h(0, u_h) - u_h$ and $y_l = \theta_l q_l^e$. We have $\frac{\partial y_l}{\partial t_l} = 0$, $\frac{\partial y_l}{\partial t_h} = 0$ and $\frac{\partial y_h}{\partial t_l} = 0$. Finally notice that

$$\frac{\partial y_h}{\partial t_h} = \theta_h \frac{\partial q_h(u_l, u_h)}{\partial \Delta} \frac{\partial \Delta}{\partial t_h} - \frac{\partial u_h}{\partial t_h} = \left[\frac{\theta_h}{\Delta \theta} - 1 \right] \frac{\partial \Delta}{\partial t_h} < 0,$$

as claimed.

Case 3 Assume that (t_l, t_h) belong to an open set $U \in D_-$.

Case 3.1 Assume that IC_h binds and IR does not in U .

The equilibrium is given by the following equations:

$$\begin{aligned} S_l(u_l, u_h) + t_l \left(\frac{\partial S_l(u_l, u_h)}{\partial u_l} - 1 \right) &= u_l \\ S_h^e - t_h + t_h \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(u_l, u_h)}{\partial u_h} &= u_h. \end{aligned}$$

Prices are given by $y_l = \theta_l q_l(u_l, u_h) - u_l$ and $y_h = \theta_h q_h^e - u_h$. Comparative statics are given by: $\frac{\partial y_h}{\partial t_l} = -\frac{\partial u_h}{\partial t_l} > 0$, $\frac{\partial y_h}{\partial t_h} = -\frac{\partial u_h}{\partial t_h} > 0$, $\frac{\partial y_l}{\partial t_l} = \theta_l \frac{\partial q_l(u_l, u_h)}{\partial \Delta} \frac{\partial \Delta}{\partial t_l} - \frac{\partial u_l}{\partial t_l} > 0$ and

$$\frac{\partial y_l}{\partial t_h} = \theta_l \frac{\partial q_l(u_l, u_h)}{\partial \Delta} \frac{\partial \Delta}{\partial t_h} - \frac{\partial u_l}{\partial t_h} = [\varphi'(q_l(u_l, u_h) \Delta \theta - \lambda t_l \varphi''(q_l(u_l, u_h))] \left(\frac{1}{\Delta \theta} \right)^2 \frac{\partial \Delta}{\partial t_h}.$$

We must show that if $\varphi''' \leq 0$ then y_l is quasi-convex in t_h . For that notice that $\text{sign} \left[\frac{\partial y_l}{\partial t_h} \right] = \text{sign} [\lambda t_l \varphi''((q_l(\Delta(t_h)) - \varphi'(q_l(\Delta(t_h)) \Delta \theta)]$. Notice that $q_l \rightarrow \lambda t_l \varphi''(q_l) - \varphi'(q_l) \Delta \theta$ is decreasing in q_l and $\frac{\partial q_l}{\partial t_h} < 0$. Therefore $\frac{\partial y_l}{\partial t_h}(t_h^*) = 0$ for some t_h^* implies $\frac{\partial y_l}{\partial t_h}(\tilde{t}_h) > 0$ for $\tilde{t}_h > t_h^*$.

Case 3.2 Assume that IC_h and IR bind and $u_h > 0$ in U .

The equilibrium is given by:

$$S_h^e - \lambda t_h + \lambda t_h \left(\frac{p_l}{p_h} \right) \frac{\partial S_l(0, u_h)}{\partial u_h} = u_h$$

We have $y_h = \theta_h q_h^e - u_h$ and $y_l = \theta_l q_l(u_l, u_h)$. We have the following comparative statics: $\frac{\partial y_l}{\partial t_l} = \frac{\theta_l}{\Delta \theta} \frac{\partial \Delta}{\partial t_l} = 0$, $\frac{\partial y_l}{\partial t_h} = \frac{\theta_l}{\Delta \theta} \frac{\partial \Delta}{\partial t_h} < 0$, $\frac{\partial y_h}{\partial t_l} = -\frac{\partial u_h}{\partial t_l} = 0$ and $\frac{\partial y_h}{\partial t_h} = -\frac{\partial u_h}{\partial t_h} > 0$.

Case 3.3 Assume that $u_h = 0$ in U .

In this case, $\frac{\partial y_k}{\partial t_l} = 0$ for $(k, j) \in \{l, h\}^2$.

This completes the proof. ■

Proof of Proposition 4. We first claim that in any (pure-strategy) equilibrium (u_l, u_h) each firm makes all profits from the low type. In fact, otherwise the firm could profitably deviate to

$(u_l + \varepsilon, u_h + \varepsilon)$ for some small $\varepsilon > 0$. Next notice that if the IC_h binds then firm could profitably deviate by offering $(u_l, u_l + \Delta\theta q_l^e)$. We conclude that IC_h does not bind and hence any equilibrium must satisfy

$$(u_l, u_h) \in \arg \max_{u_l^i, u_h^i} \left\{ p_l H(u_l^i - u_l) (S_l^e - u_l^i) + p_h \hat{D}_h^i(u_h^i | u_h) (S_h(u_l^i, u_h^i) - u_h^i) \right\},$$

where

$$\hat{D}_h^i(u_h^i | u_h) := \begin{cases} 1 & \text{if } u_h^i > u_h \\ 0 & \text{if } u_h^i < u_h \\ \frac{1}{J} & \text{if } u_h^i = u_h. \end{cases}$$

Since each firm makes zero profits from high types, $u_h = S_h(u_l, u_h)$. Therefore, we must also have

$$u_l \in \arg \max_{u_l^i} \left\{ p_l H(u_l^i - u_l) (S_l^e - u_l^i) \right\},$$

which allows us to conclude that $u_l = \max\{S_l^e - \lambda t_l, 0\}$. Hence u_h solves

$$S_h(\max\{S_l^e - \lambda t_l, 0\}, u_h) - u_h = 0.$$

Recall that $\Delta(u_l, u_h) := u_h - u_l$ and hence the equilibrium requires that $S_h(0, \Delta(u_l, u_h)) - u_h = 0$. Observe that

$$\left(\frac{\partial}{\partial u_h} \right) [S_h(0, \Delta(u_l, u_h)) - u_h] < 0.$$

This implies that whenever

$$\Delta(\max\{S_l^e - t_l, 0\}, S_h^*) \leq \Delta\theta q_h^e,$$

any pure-strategy equilibrium should be $(\max\{S_l^e - \lambda t_l, 0\}, S_h^*)$, while, if

$$\Delta(\max\{S_l^e - \lambda t_l, 0\}, S_h^e) > \Delta\theta q_h^e,$$

then any pure-strategy equilibrium should be

$$(\max\{S_l^e - \lambda t_l, 0\}, \hat{u}_h(\max\{S_l^e - \lambda t_l, 0\})),$$

where $\hat{u}_h(u_l)$ satisfies $S_h(u_l, \hat{u}_h(u_l)) - \hat{u}_h(u_l) = 0$.

Next assume that $0 \leq t_l \leq \frac{S_l^e - \eta_h}{\lambda}$, and notice that this implies that $\max\{S_l^e - \lambda t_l, 0\} = S_l^e - \lambda t_l$. Moreover, in this case the problem is separable in the utility of each type, which easily implies that $u_l^* = S_l^e - \lambda t_l$ and $u_h^* = S_h^e$ is an equilibrium.

Next assume that $\frac{S_l^e - \eta_h}{\lambda} < t_l < \tilde{t}_l$. Recall that the unique putative pure-strategy equilibrium is

$$(u_l^*, u_h^*) = (\max\{S_l^e - \lambda t_l, 0\}, \hat{u}_h(\max\{S_l^e - \lambda t_l, 0\})).$$

We will show that the firm has a profitable deviation with the following structure: the firm relinquishes an $\varepsilon \approx 0$ more utility to high types (conquering the entire type- h market), and chooses u_l to solve

$$\sup_{u_l > \max\{S_l^e - \lambda t_l, 0\}} \left\{ p_l H \left(\frac{u_l - u_l^*}{t_l} \right) (S_l^e - u_l) + p_h (S_h(u_l, \hat{u}_h(u_l^*)) - u_h^*) \right\}.$$

The right-derivative at $u_l = u_l^*$ is

$$p_l h(0) \left(\frac{S_l^e - u_l^*}{t_l} \right) - p_l H(0) + p_h \frac{\partial S_h(u_l^*, \hat{u}_h(u_l^*))}{\partial u_l}.$$

Since $H(0) = \frac{1}{j}$, the last expression is proportional to

$$(S_l^e - u_l^*) - \lambda t_l + \lambda t_l J \frac{p_h}{p_l} \frac{\partial S_h(u_l^*, \hat{u}_h(u_l^*))}{\partial u_l}.$$

If $u_l^* > 0$ then $p_l (S_l^e - u_l^*) - \lambda t_l = 0$ and hence the expression above is

$$\lambda t_l J \frac{p_h}{p_l} \frac{\partial S_h(u_l^*, \hat{u}_h(u_l^*))}{\partial u_l} > 0.$$

Otherwise the expression above is

$$S_l^e - \lambda t_l + \lambda t_l J \frac{p_h}{p_l} \frac{\partial S_h(0, \hat{u}_h)}{\partial u_l} > 0,$$

because $t_l < \tilde{t}_l$.

Finally, assume that $t_l \geq \tilde{t}_l$. We will show that $u_l^* = 0$ and \hat{u}_h is a pure-strategy equilibrium. The argument in the beginning of this proof then implies that this is the unique pure-strategy equilibrium. We start arguing that the firm cannot improve by offering a menu in which the utility of the high type is smaller than \hat{u}_h . This is because any such menu would attract no high types, while the firm would not profit from low types as the menu $(0, \hat{u}_h)$ maximizes the firm's utility when it considers only low types. Moreover, since the firm obtains zero profits from high types, a similar argument establishes that the firm cannot increase its profits by offering a menu that gives utility \hat{u}_h to high types. Therefore it remains to argue that the firm cannot profit by offering menus of the kind $(\delta_l, \hat{u}_h + \delta_h)$, where $(\delta_l, \delta_h) \in \mathbb{R}_+ \times \mathbb{R}_{++}$. In light of the previous argument, we consider the following upper bound to the firm's utility in which the constraint IC_h is ignored and it is assumed that the firm attracts all high types whenever such consumers yield positive profits:

$$G(\delta_l, \delta_h) \equiv p_l H \left(\frac{\delta_l}{t_l} \right) (S_l^e - \delta_l) + p_h J \max \{ S_h(\delta_l, \hat{u}_h + \Delta u_h) - \hat{u}_h - \Delta u_h, 0 \}.$$

Therefore, if we show that

$$(0, 0) \in \operatorname{argmax}_{(\delta_l, \Delta u_h) \in \mathbb{R}_+ \times \mathbb{R}_{++}} G(\delta_l, \Delta u_h),$$

then we will have concluded that $(0, \hat{u}_h)$ is an equilibrium. First notice that any deviation $(\delta_l, \delta_h) \in \mathbb{R}_+ \times \mathbb{R}_{++}$ in which $\delta_h \geq \delta_l$ is weakly dominated by $(\delta_l, \delta_h) = (0, 0)$ because it decreases the profits

obtained from both types. Henceforth we can restrict attention to deviations $(\delta_l, \delta_h) \in \mathbb{R}_+ \times \mathbb{R}_{++}$ in which $\delta_h < \delta_l$. Again, since $\delta_h \rightarrow S_h(\delta_l, \hat{u}_h + \delta_h) - \hat{u}_h - \delta_h$ is strictly decreasing in δ_h , a sufficient condition for the absence of profitable deviation is that $G(0, 0) \geq G(\delta_l, 0)$ for all $\delta_l > 0$. First, take $\delta_l > 0$, notice that $S_h(\delta_l, \hat{u}_h) > S_h(0, \hat{u}_h)$. Next we claim that $\delta_l \rightarrow G(\delta_l, 0)$ is quasi-concave. Indeed take $\delta_l > 0$ and assume that

$$0 \geq G_1(\delta_l, 0) = p_l h \left(\frac{\delta_l}{t_l} \right) \left(\frac{S_l^e - \delta_l}{t_l} \right) - p_l H \left(\frac{\delta_l}{t_l} \right) + p_h \frac{\partial S_h(\delta_l, \hat{u}_h)}{\partial \delta_l} \\ p_l H \left(\frac{\delta_l}{t_l} \right) \left[\frac{h \left(\frac{\delta_l}{t_l} \right)}{H \left(\frac{\delta_l}{t_l} \right)} \left(\frac{S_l^e - \delta_l}{t_l} \right) - 1 \right] + p_h \frac{\partial S_h(\delta_l, \hat{u}_h)}{\partial \delta_l}.$$

Notice that $\frac{\partial^2 S_h(\delta_l, \hat{u}_h)}{\partial \delta_l^2} < 0$ and, since H is strictly logconcave, $\delta_l \rightarrow \frac{h \left(\frac{\delta_l}{t_l} \right)}{H \left(\frac{\delta_l}{t_l} \right)}$ is decreasing in δ_l .

Therefore $\tilde{\delta}_l > \delta_l$ implies $G_1(\tilde{\delta}_l, 0) \leq 0$. Therefore, it suffices to show that $G_{1+}(0, 0) \leq 0$, which holds because $G_{1+}(0, 0)$ has the same sign as

$$\frac{1}{J} \frac{p_l}{p_h} \left(\frac{S_l^e}{\lambda t_l} - 1 \right) + \frac{\partial S_h}{\partial u_l}(0, \hat{u}_h) \leq 0,$$

because $t_l \geq \tilde{t}_l$. ■

Proof of Proposition 5. Let $k = l, h$. Consider a (symmetric) mixed strategy equilibrium F , with marginals F_l and F_k . Write F^{-i} for the distribution over utilities when every firm $j \neq i$ plays F . Write F_k^{-i} for its marginal. Let \underline{u}_k be the infimum over all utilities provided to type k and let \bar{u}_k be the supremum. Since the profit function is supermodular in (u_l, u_k) , we conclude that $(\underline{u}_l, \underline{u}_h)$ and (\bar{u}_l, \bar{u}_h) are optimal. We first prove four lemmas.

Lemma 1. *For any mixed-strategy equilibrium, if the constraint IC_k ($k \in \{l, h\}$) is slack at some optimal menu then it is slack at any menu in the support of F .*

Proof of Lemma 1. Write $-k$ for $j \in \{l, h\}$, $j \neq k$. Assume that IC_{-k} is slack at $(\underline{u}_l, \underline{u}_h)$. Then

$$\frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)] \leq 0. \quad (19)$$

Assume towards a contradiction that there exists an optimal menu (u_l, u_h) at which IC_{-k} is not slack. Then $u_k^i > \underline{u}_k$ and hence since the payoff is differentiable

$$0 = \frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k(u_l, u_h) - u_k^i)] \\ = \frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)] + \frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e(u_l, u_h) - S_k^e)] \\ < \frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)] \leq 0,$$

where the first inequality used $(S_k^e(u_l, u_h) - S_k^e) < 0$ and $\frac{\partial}{\partial u_k^i} S_k^e(u_l, u_h) < 0$ and the second used (19) and the assumption that $u_k^i \rightarrow \mathbb{E}_{F_k^{-i}} [D(u_k^i, u_k^{-i}) (S_k^e - u_k^i)]$ is strictly quasi-concave.

Next assume that IC_k binds at $(\underline{u}_l, \underline{u}_h)$. Notice that this implies that IC_{-k} is slack at $(\underline{u}_l, \underline{u}_h)$ and hence at every optimal menu. Moreover, this implies

$$\frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)] < 0.$$

Hence assume towards a contradiction that there exists (u_l, u_h) at which IC_k is slack. Then, since $u_k^i > \underline{u}_k$ and since IC_{-k} is also slack at $(\underline{u}_l, \underline{u}_h)$, we have

$$\frac{\partial}{\partial u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)] = 0,$$

which is a contradiction because $u_k^i \rightarrow \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)]$ is strictly quasi-concave. \blacksquare

Lemma 2. *For any (symmetric) equilibrium, the constraint IC_l binds only if $\Lambda(t_l, t_h) < \eta_h$.*

Proof of Lemma 2. Assume that IC_l binds for some menu in the support of F . By Lemma 1, it binds for every menu. Hence,

$$0 \geq \frac{\partial}{\partial u_l^i} \mathbb{E}_{F_l^{-i}} [D_l^i(\underline{u}_l, u_l^{-i}) (S_l^e - \underline{u}_l)] = \mathbb{E}_{F_l^{-i}} \left[D_l^i(\underline{u}_l, u_l^{-i}) \left[\frac{\frac{\partial D_l^i(\underline{u}_l, u_l^{-i})}{\partial u_l^i}}{D_l^i(\underline{u}_l, u_l^{-i})} (S_l^e - \underline{u}_l) - 1 \right] \right].$$

By Assumption 3, $\frac{\partial D_l^i(\underline{u}_l, u_l^{-i})}{\partial u_l^i}$ is increasing in u_l^{-i} . Therefore

$$0 \geq \mathbb{E}_{F_l^{-i}} \left[D_l^i(\underline{u}_l, u_l^{-i}) \left[\frac{\frac{\partial D_l^i(\underline{u}_l, u_l^{-i})}{\partial u_l^i}}{D_l^i(\underline{u}_l, u_l^{-i})} (S_l^e - \underline{u}_l) - 1 \right] \right],$$

which implies

$$\underline{u}_l \geq S_l^e - \lambda t_l. \quad (20)$$

Similarly, optimality for u_h at (\bar{u}_l, \bar{u}_h) implies

$$\begin{aligned} 0 &< \frac{\partial}{\partial u_h^i} \mathbb{E}_{F_h^{-i}} [D_h^i(\bar{u}_h, u_h^{-i}) (S_h(\bar{u}_l, \bar{u}_h) - \bar{u}_h)] \\ &= \frac{\partial}{\partial u_h^i} \mathbb{E}_{F_h^{-i}} [D_h^i(\bar{u}_h, u_h^{-i}) (S_h(\bar{u}_l, \bar{u}_h) - S_h^e)] + \mathbb{E}_{F_h^{-i}} \left[D_h^i(\bar{u}_h, u_h^{-i}) \left[\frac{\frac{\partial D_h^i(\bar{u}_h, u_h^{-i})}{\partial u_h^i}}{D_h^i(\bar{u}_h, u_h^{-i})} (S_h^e - \bar{u}_h) - 1 \right] \right], \end{aligned}$$

which implies

$$0 < \mathbb{E}_{F_h^{-i}} \left[D_h^i(\bar{u}_h, u_h^{-i}) \left[\frac{\frac{\partial D_h^i(\bar{u}_h, u_h^{-i})}{\partial u_h^i}}{D_h^i(\bar{u}_h, u_h^{-i})} (S_h^e - \bar{u}_h) - 1 \right] \right],$$

and, using the fact that $\frac{\partial D_h^i(\bar{u}_h, u_h^{-i})}{D_h^i(\bar{u}_h, u_h^{-i})}$ is increasing by Assumption 3,

$$< \mathbb{E}_{F_h^{-i}} \left[D_h^i(\bar{u}_h, u_h^{-i}) \left[\frac{\partial D_h^i(\bar{u}_h, \bar{u}_h^{-i})}{D_h^i(\bar{u}_h, \bar{u}_h^{-i})} (S_h^e - \bar{u}_h) - 1 \right] \right],$$

implying that

$$\bar{u}_h < S_h^e - \lambda t_h. \quad (21)$$

Therefore, using (20) and (21) we have

$$\begin{aligned} \Lambda(t_l, t_h) &\equiv \lambda t_h + \max\{S_l^e - \lambda t_l, 0\} \\ &< S_h^e - (\bar{u}_h - \underline{u}_l) \\ &< S_h^e - (\bar{u}_h - \bar{u}_l) \\ &< S_h^e - \Delta\theta q_h^e, \end{aligned}$$

which concludes the proof. ■

Lemma 3. *For any (symmetric) equilibrium, the constraint IC_h binds only if $\Lambda(t_l, t_h) > \bar{\eta}$.*

Proof of Lemma 3. Using the optimality condition for \bar{u}_h and an argument similar to the one from Lemma 2, we conclude that, if IC_h binds, then

$$\mathbb{E}_{F_h^{-i}} \left[D_h^i(\underline{u}_h, u_h^{-i}) \left[\frac{\partial D_h^i(\underline{u}_h, u_h^{-i})}{D_h^i(\underline{u}_h, u_h^{-i})} (S_h^e - \underline{u}_h) - 1 \right] \right] < 0.$$

Since $\frac{\partial D_h^i(\underline{u}_h, u_h^{-i})}{D_h^i(\underline{u}_h, u_h^{-i})}$ is decreasing in u_h^{-i} , we obtain that

$$\mathbb{E}_{F_h^{-i}} \left[D_h^i(\underline{u}_h, \underline{u}_h^{-i}) \left[\frac{\partial D_h^i(\underline{u}_h, \underline{u}_h^{-i})}{D_h^i(\underline{u}_h, \underline{u}_h^{-i})} (S_h^e - \underline{u}_h) - 1 \right] \right] < 0$$

and $\underline{u}_h > S_h^e - \lambda t_h$.

Similarly, we get that

$$\bar{u}_l \leq \max\{S_l^e - \lambda t_l, 0\}.$$

Therefore,

$$\begin{aligned} \Lambda(t_l, t_h) &\equiv \lambda t_h + \max\{S_l^e - \lambda t_l, 0\} \\ &> S_h^e - (\underline{u}_h - \bar{u}_l) \\ &\geq S_h^e - (\underline{u}_h - \underline{u}_l) \\ &> S_h^e - \Delta\theta q_l^e, \end{aligned}$$

which concludes the proof. ■

Lemma 4. For any (symmetric) equilibrium, if no IC constraint binds then $\Lambda(t_l, t_h) \in [\eta_h, \bar{\eta}]$.

Proof of Lemma 4. If no IC constraint binds then if $(\tilde{u}_l, \tilde{u}_h)$ is optimal then for $k = l, h$ we have

$$\tilde{u}_k \in \operatorname{argmax}_{u_k^i} \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)].$$

Since by assumption $u_k^i \rightarrow \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)]$ is strictly quasi-concave, the equilibrium is in pure strategies. Hence if (u_l^*, u_h^*) is the pure-strategy equilibrium, using the optimality conditions and symmetry one gets

$$u_l^* = \max \{S_l^e - \lambda t_l, 0\}$$

and

$$u_h^* = S_h^e - \lambda t_h > 0,$$

where the last inequality used the fact that no IC constraint binds by assumption. Therefore

$$\begin{aligned} (t_l, t_h) &\equiv \lambda t_h + \max \{S_l^e - \lambda t_l, 0\} \\ &= S_h^e - (u_h^* - u_l^*) \in [\eta_h, \bar{\eta}], \end{aligned}$$

where the last conclusion follows from the assumption that no IC constraint binds. \blacksquare

We are now ready to prove the Proposition.

Claim (a). Assume that $(t_l, t_h) \in E$. By Lemmas 1-4, no IC constraint binds. Hence since $u_k^i \rightarrow \mathbb{E}_{F_k^{-i}} [D_k^i(u_k^i, u_k^{-i}) (S_k^e - u_k^i)]$ is strictly quasi-concave, the equilibrium is in pure strategies.

Claim (b). Assume that $(t_l, t_h) \in D_+$. By Lemmas 1-4, the constraint IC_l binds in every equilibrium. Consider a menu (u_l, u_h) in which $u_l > 0$ and $\mathbb{E}_{F_l^{-i}} [D_h^i(u_h, u_h^{-i})] > 0$ (menus satisfying $\mathbb{E}_{F_l^{-i}} [D_h^i(u_h, u_h^{-i})] = 0$ are obviously suboptimal) and notice that the partial derivative of the profit function yields

$$\mathbb{E}_{F_l^{-i}} \left[\frac{\partial D_h^i(u_h, u_h^{-i})}{\partial u_h} \left(\frac{\partial S_h(u_l, u_h)}{\partial u_l} \right) \right] + \mathbb{E}_{F_l^{-i}} \left[D_h^i(u_h, u_h^{-i}) \left(\frac{\partial S_h^2(u_l, u_h)}{\partial u_l \partial u_h} \right) \right] > 0.$$

Because $\frac{\partial S_h(u_l, u_h)}{\partial u_l} > 0$ and $\left(\frac{\partial S_h^2(u_l, u_h)}{\partial u_l \partial u_h} \right) > 0$, this implies that the profit function is strictly super-modular and hence every equilibrium is ordered.

Next we fix $t_l \in (S_l^e - \eta_h, \tilde{t}_l)$ and show that there exists $\hat{\tau}_h(t_l) > 0$ (with $(t_l, \hat{\tau}_h(t_l)) \in D_+$) such that all equilibria are in mixed strategies whenever $t_h < \hat{\tau}_h(t_l)$. For that we will prove that there does not exist a pure-strategy equilibrium whenever t_h is small. Indeed, parametrizing equilibria by t_h , any putative pure-strategy equilibrium $(u_l^*(t_h), u_h^*(t_h))$ satisfies

$$S_h(u_l^*(t_h), u_h^*(t_h)) - u_h^*(t_h) + \lambda t_h \left(\frac{\partial S_h}{\partial u_h}(u_l^*(t_h), u_h^*(t_h)) - 1 \right) = 0 \quad (22)$$

and

$$(S_l^e - u_l^*(t_h)) + \lambda t_l \left(\frac{p_h}{p_l} \frac{\partial S_h}{\partial u_l}(u_l^*(t_h), u_h^*(t_h)) - 1 \right) \leq 0, \quad (23)$$

with equality if $u_l^*(t_h) > 0$. From (22) we see that $S_h(u_l^*(t_h), u_h^*(t_h)) - u_h^*(t_h) \rightarrow 0$ as $t_h \rightarrow 0$. Using this and (23), we see that

$$\lim_{t_h \downarrow 0} (u_l^*(t_h), u_h^*(t_h)) = (u_l^*(0), u_h^*(0)),$$

where $(u_l^*(0), u_h^*(0))$ is the pair that solves (22) and (23) when $t_h = 0$. Recall from Proposition 4 that $(u_l^*(0), u_h^*(0))$ is not an equilibrium as a firm would have a deviation to some $(u_l^\dagger, u_h^\dagger)$ with $u_l^\dagger > u_l^*(0)$ and $u_h^\dagger > u_h^*(0)$. Writing $\Pi((u_l, u_h) | (u_l^*(t_h), u_h^*(t_h)), t_h)$ for the profit from (u_l, u_h) when all other firms play $(u_l^*(t_h), u_h^*(t_h))$ and parameters are (t_l, t_h) , we conclude that there exists $\varepsilon > 0$ such that

$$\Pi\left(\left(u_l^\dagger, u_h^\dagger\right) \mid \left(u_l^*(0), u_h^*(0)\right), 0\right) > \Pi\left(\left(u_l^*(0), u_h^*(0)\right) \mid \left(u_l^*(0), u_h^*(0)\right), 0\right) + \varepsilon.$$

Since

$$\lim_{t_h \downarrow 0} \Pi\left(\left(u_l^*(t_h), u_h^*(t_h)\right) \mid \left(u_l^*(t_h), u_h^*(t_h)\right), t_h\right) = \Pi\left(\left(u_l^*(0), u_h^*(0)\right) \mid \left(u_l^*(0), u_h^*(0)\right), 0\right)$$

and

$$\lim_{t_h \downarrow 0} \Pi\left(\left(u_l^\dagger, u_h^\dagger\right) \mid \left(u_l^*(t_h), u_h^*(t_h)\right), t_h\right) = \Pi\left(\left(u_l^\dagger, u_h^\dagger\right) \mid \left(u_l^*(0), u_h^*(0)\right), 0\right),$$

we conclude that there exists $\hat{\tau}_h(t_l) > 0$ such that

$$\Pi\left(\left(u_l^\dagger, u_h^\dagger\right) \mid \left(u_l^*(t_h), u_h^*(t_h)\right), t_h\right) > \Pi\left(\left(u_l^*(t_h), u_h^*(t_h)\right) \mid \left(u_l^*(t_h), u_h^*(t_h)\right), t_h\right)$$

for every $t_h < \hat{\tau}_h(t_l)$, which proves the claim.

Claim (c). The proof is similar to the one from (b) and therefore omitted. ■

Proofs of Proposition 6 and Corollary 2. It is slightly more convenient to perform a change of variables so that $t(\theta) = \alpha + \beta\theta > 0$, where $\alpha > 0$. To see the equivalence of the formulations take $\alpha = \underline{t} - (\bar{t} - \underline{t}) \left(\frac{\theta}{\bar{\theta} - \underline{\theta}}\right)$ and $\beta = \left(\frac{\bar{t} - \underline{t}}{\bar{\theta} - \underline{\theta}}\right)$.

Existence of Symmetric Equilibrium when $|\Delta t|$ is small. In the new parametrization, we must show existence when $|\beta|$ is small. It is straightforward to verify that the firm's problem is equal to the problem of Roche and Stole (2002) when $\beta = 0$ and α is small. It follows from Roche and Stole (2002) that there exists $\zeta_1 > 0$ such that if $t(\theta)$ is a constant belonging to $(0, \zeta_1)$ then the problem has a unique solution in which $\dot{u}(\theta) = \theta$ for every θ . Consider $t(\theta) = \alpha + \beta\theta$. A straightforward continuity argument implies that for every ζ in this range there exists $\varepsilon_1 > 0$ such that $\sup_{\theta \in [\underline{\theta}, \bar{\theta}]} \|\zeta - \alpha + \beta\theta\| < \varepsilon_1$ then if $(u_{\alpha, \beta}(\theta))_{\theta \in [\underline{\theta}, \bar{\theta}]}$ satisfies the Euler equation above with the same boundary conditions (which solution is continuous in α, β) then if $(\hat{u}_{\alpha, \beta}(\theta))_{\theta \in [\underline{\theta}, \bar{\theta}]}$ is a best-response to this equation we must have $\left(\frac{1}{2} + \frac{\hat{u}_{\alpha, \beta}(\theta) - u_{\alpha, \beta}(\theta)}{2t(\theta)}\right) \in \left[\frac{1}{3}, \frac{2}{3}\right]$ for every θ . Without loss of generality, we make this restriction for the remainder of this proof. In light of this, we set up the

optimal control problem where $y(\theta)$ is the control variable, $x(\theta)$ is the state variable and the costate variable is $p(\theta)$:

$$H(\theta, x(\theta), y(\theta), p(\theta)) = \left(\frac{1}{2} + \frac{x(\theta) - \tilde{u}(\theta)}{2t(\theta)} \right) \left(\theta y(\theta) - \frac{(y(\theta))^2}{2} - x(\theta) \right) + p(\theta) y(\theta).$$

Optimality Conditions are $H_y(\theta, x(\theta), y(\theta), p(\theta)) = 0, -\dot{p}(\theta) = H_x(\theta, x(\theta), y(\theta), p(\theta)), \dot{x}(\theta) = y(\theta)$ and the transversality conditions $y(\bar{\theta}) = \bar{\theta}, y(\underline{\theta}) = \underline{\theta}$. Solving and imposing symmetry we obtain

$$\begin{aligned} \left(\frac{1}{2} \right) (\theta - y(\theta)) &= -p(\theta) \\ -\dot{p}(\theta) &= \left(\frac{1}{2t(\theta)} \right) \left(\theta y(\theta) - \frac{(y(\theta))^2}{2} - x(\theta) \right) - \left(\frac{1}{2} \right). \end{aligned}$$

Let $(x_{\alpha,\beta}^*(\theta), y_{\alpha,\beta}^*(\theta), p_{\alpha,\beta}^*(\theta))$ be a solution to the system above. We claim that this is an optimal solution for the firm's problem. To show this we will use Arrow sufficiency condition to show that it is optimal for each firm when we ignore monotonicity constraints and IR constraints. We will then verify that the solution satisfies such constraints. Define:

$$H(\theta, x(\theta), p_{\alpha,\beta}^*(\theta)) = \max_{\tilde{y}} \left(\frac{1}{2} + \frac{x(\theta) - \hat{u}_{\alpha,\beta}(\theta)}{2t(\theta)} \right) \left(\theta \tilde{y} - \frac{(\tilde{y})^2}{2} - x(\theta) \right) + p_{\alpha,\beta}^*(\theta) \tilde{y}$$

Letting $A(x) := \left(\frac{1}{2} + \frac{x(\theta) - \hat{u}_{\alpha,\beta}(\theta)}{2t(\theta)} \right)$ and $B(x) := A(x)\theta + p_{\alpha,\beta}^*(\theta)$, we obtain

$$H(\theta, x(\theta), p^*(\theta)) = \frac{B(x)^2}{2A(x)} - x(\theta)A(x).$$

Arrow sufficiency condition will be satisfied if we guarantee that $H_{xx}(\theta, x(\theta), p^*(\theta)) \leq 0$. Since $x \rightarrow A(x) = \left(\frac{1}{2} + \frac{x(\theta) - \hat{u}_{\alpha,\beta}(\theta)}{2t(\theta)} \right)$ is a positive affine transformation, we may change variables and get $A(x) = z$, $B(x) = z\theta + p_{\alpha,\beta}^*(\theta)$ and $x(\theta) = 2zt(\theta) - t(\theta) + \hat{u}_{\alpha,\beta}(\theta)$, which implies

$$H(\theta, z, p^*(\theta)) = \frac{(z\theta + p_{\alpha,\beta}^*)^2}{2z} - (2zt(\theta) - t(\theta) + \hat{u}_{\alpha,\beta}(\theta))z.$$

Differentiating twice with respect to z we get

$$\begin{aligned} H_{zz}(\theta, z, p^*(\theta)) &= \frac{p_{\alpha,\beta}^{*2}}{z^3} - 4t(\theta) \\ &= \frac{(\theta - y_{\alpha,\beta}^*(\theta))^2}{4z^3} - 4t(\theta) \\ &< \frac{9(\theta - \hat{u}_{\alpha,\beta}(\theta))^2}{4} - 4(\alpha - |\beta|\theta), \end{aligned}$$

where the last line used $z = \left(\frac{1}{2} + \frac{\hat{u}_{\alpha,\beta}(\theta) - u_{\alpha,\beta}(\theta)}{2t(\theta)} \right) \in \left[\frac{1}{3}, \frac{2}{3} \right]$ and $y_{\alpha,\beta}^*(\theta) = \hat{u}_{\alpha,\beta}(\theta)$. The solution $(\hat{u}_{\alpha,\beta}(\theta))$ when $\beta = 0$ satisfies $(\theta - \hat{u}_{\alpha,0}(\theta)) = 0$. By a continuity argument, there exists $\varepsilon_2 \in (0, \varepsilon_1)$

such that $\sup_{\theta \in [\underline{\theta}, \bar{\theta}]} \|\zeta - \alpha + \beta\theta\| < \varepsilon_2$ implies $\max_{\theta} \frac{9(\theta - \hat{u}_{\alpha, \beta}(\theta))^2}{4} - 4(\alpha - |\beta|\theta) < 0$, which guarantees optimality of the relaxed program that ignores monotonicity and IR constraints. Moreover, a continuity argument implies that we can take $\varepsilon \in (0, \varepsilon_2)$ such that $\sup_{\theta \in [\underline{\theta}, \bar{\theta}]} \|\zeta - \alpha + \beta\theta\| < \varepsilon$ implies that $\hat{u}_{\alpha, \beta}(\theta)$ is increasing and $\hat{u}_{\alpha, \beta}(\underline{\theta}) > 0$. This completes the proof.

Distortions when $\Delta t > 0$ and when $\Delta t < 0$. Notice that according to our change of variables β and $\Delta t > 0$ have the same sign.

Case 1: $\beta > 0$. Let $(u(\theta))_{\theta \in [\underline{\theta}, \bar{\theta}]}$ be the solution when $t(\theta) = \alpha + \beta\theta$. We will show that $\dot{u}(\theta) < \theta$ for all $\theta \in (\underline{\theta}, \bar{\theta})$.

First we show that $\ddot{u}(\underline{\theta}) < 1$. Assume towards a contradiction that $\ddot{u}(\underline{\theta}) \geq 1$. Let $(u_{\lambda}(\theta))_{\theta \in [\underline{\theta}, \bar{\theta}]}$ be the solution when $t(\theta) = \alpha + \beta\underline{\theta}$ for every θ and notice that $\ddot{u}(\theta) = 1$ for every θ . Notice that

$$\begin{aligned} \frac{d\ddot{u}(\theta)}{d\theta} \Big|_{\underline{\theta}} &= - \left(\frac{1}{(\alpha + \beta\underline{\theta})^2} \right) [(\underline{\theta} - \dot{u}(\underline{\theta})) \ddot{u}(\theta) + (\underline{\theta} - \dot{u}(\underline{\theta}))] (\alpha + \beta\underline{\theta})^2 + \beta \left(\frac{1}{(\alpha + \beta\underline{\theta})^2} \right) \left(\theta \dot{u}(\theta) - \frac{(\dot{u}(\theta))^2}{2} - u(\theta) \right) \\ &= \beta \left(\frac{1}{(\alpha + \beta\underline{\theta})^2} \right) \left(\theta \dot{u}(\theta) - \frac{(\dot{u}(\theta))^2}{2} - u(\theta) \right) > 0. \end{aligned}$$

Therefore, there exists $\varepsilon > 0$ such that, for all $\theta \in (\underline{\theta}, \underline{\theta} + \varepsilon)$ we have $\ddot{u}(\theta) > 1$, implying

$$\theta \dot{u}(\theta) - \frac{(\dot{u}(\theta))^2}{2} - u(\theta) < \theta \dot{u}_{\lambda}(\theta) - \frac{(\dot{u}_{\lambda}(\theta))^2}{2} - u_{\lambda}(\theta).$$

Therefore $\dot{u}(\theta) > \theta$ for every θ in this set. Let θ^* be the infimum over all $\theta \geq \underline{\theta} + \varepsilon$ such that $\dot{u}(\theta) \leq \theta$. Since $\dot{u}(\bar{\theta}) = \bar{\theta}$, we conclude that θ^* is well defined. Notice that since $\dot{u}(\theta) > \theta$ for every $\theta \in (\underline{\theta}, \theta^*)$ we have $\ddot{u}(\theta^*) \leq \ddot{u}_{\lambda}(\theta^*) = 1$. However, since $u(\theta^*) > u_{\lambda}(\theta^*)$ we must have

$$\begin{aligned} \ddot{u}(\theta^*) &= 2 - \left(\frac{1}{\alpha + \beta\theta^*} \right) \left(\theta^* \dot{u}(\theta^*) - \frac{(\dot{u}(\theta^*))^2}{2} - u(\theta^*) \right) \\ &> 2 - \left(\frac{1}{\alpha + \beta\underline{\theta}} \right) \left(\theta^* \dot{u}(\theta^*) - \frac{(\dot{u}(\theta^*))^2}{2} - u(\theta^*) \right) \\ &> 2 - \left(\frac{1}{\alpha + \beta\underline{\theta}} \right) \left(\theta^* \dot{u}(\theta^*) - \frac{(\dot{u}(\theta^*))^2}{2} - u_{\lambda}(\theta^*) \right) \\ &= \ddot{u}_{\lambda}(\theta^*), \end{aligned}$$

which is a contradiction. This shows that $\ddot{u}(\underline{\theta}) < 1$ and hence there exists $\varepsilon > 0$ such that, for all $\theta \in (\underline{\theta}, \underline{\theta} + \varepsilon)$ we have $\dot{u}(\theta) < \theta$. Assume towards a contradiction that there exists $\tilde{\theta} < \bar{\theta}$ such that $\dot{u}(\tilde{\theta}) = \tilde{\theta}$. Let $\tilde{\theta}$ be the smallest type satisfying this condition. This implies $\ddot{u}(\tilde{\theta}) \geq 1$, hence the case above applies and leads to a contradiction.

Case 2: $\beta < 0$. An argument analogous to the one presented in Case 1 above establishes that $\dot{u}(\theta) > \theta$ for every $\theta \in (\underline{\theta}, \bar{\theta})$.

Types that benefit and types that are hurt by asymmetric information.

Case 1: $\beta > 0$. As we have shown above, $\ddot{u}(\underline{\theta}) < 1$, and hence

$$\begin{aligned} u(\underline{\theta}) &= (\alpha + \beta\underline{\theta}) (\ddot{u}(\underline{\theta}) - 2) + \underline{\theta} \dot{u}(\underline{\theta}) - \frac{(\dot{u}(\underline{\theta}))^2}{2} \\ &< (\alpha + \beta\underline{\theta}) (\ddot{u}_{\lambda}(\underline{\theta}) - 2) + \underline{\theta} \dot{u}_{\lambda}(\underline{\theta}) - \frac{(\dot{u}_{\lambda}(\underline{\theta}))^2}{2} \\ &= u_{\lambda}(\underline{\theta}), \end{aligned}$$

where $(u_\lambda(\theta))_{\theta \in [\underline{\theta}, \bar{\theta}]}$ is the solution when $t(\theta) = \alpha + \beta\underline{\theta}$ for every θ . This implies that type $\underline{\theta} < \theta_1$ is worse off under asymmetric information. By continuity there is θ_1 with $\underline{\theta} < \theta_1$ such that every type $\theta < \theta_1$ is worse off when there is asymmetric information. A symmetric argument shows the existence of $\theta_2 \in [\theta_1, \bar{\theta})$ such that every type $\theta > \theta_2$ is better off under asymmetric information.

Case 2: $\beta < 0$. The argument is similar to case 1 above and omitted for brevity.

An increase in β decreases $\dot{u}(\theta)$ for every interior type.

Case 1: $\beta > 0$. Keep α constant and take two solutions $u_{\beta_1}(\theta)$ and $u_{\beta_2}(\theta)$ for $\beta_2 > \beta_1$. We claim that $\dot{u}_{\beta_1}(\theta) > \dot{u}_{\beta_2}(\theta)$ for every $\theta \in (\underline{\theta}, \bar{\theta})$. First we claim that $\ddot{u}_{\beta_1}(\underline{\theta}) > \ddot{u}_{\beta_2}(\underline{\theta})$. Assume towards a contradiction that $\ddot{u}_{\beta_1}(\underline{\theta}) \leq \ddot{u}_{\beta_2}(\underline{\theta})$. Notice that $\ddot{u}_{\beta_1}(\underline{\theta}) = \ddot{u}_{\beta_2}(\underline{\theta})$ implies

$$\begin{aligned} & \frac{d\ddot{u}_{\beta_2}(\underline{\theta})}{d\theta} - \frac{d\ddot{u}_{\beta_1}(\underline{\theta})}{d\theta} \\ &= \beta_2 \left(\frac{1}{(\alpha + \beta_2 \underline{\theta})^2} \right) \left(\underline{\theta} \dot{u}_{\beta_2}(\underline{\theta}) - \frac{(\dot{u}_{\beta_2}(\underline{\theta}))^2}{2} - u_{\beta_2}(\underline{\theta}) \right) - \beta_1 \left(\frac{1}{(\alpha + \beta_1 \underline{\theta})^2} \right) \left(\underline{\theta} \dot{u}_{\beta_1}(\underline{\theta}) - \frac{(\dot{u}_{\beta_1}(\underline{\theta}))^2}{2} - u_{\beta_1}(\underline{\theta}) \right) \\ &= \left(\frac{\beta_2 \underline{\theta}}{\alpha + \beta_2 \underline{\theta}} \right) \underline{\theta}^{-1} \left(\frac{\underline{\theta} \dot{u}_{\beta_2}(\underline{\theta}) - \frac{(\dot{u}_{\beta_2}(\underline{\theta}))^2}{2} - u_{\beta_2}(\underline{\theta})}{\alpha + \beta_2 \underline{\theta}} \right) - \left(\frac{\beta_1 \underline{\theta}}{\alpha + \beta_1 \underline{\theta}} \right) \underline{\theta}^{-1} \left(\frac{\underline{\theta} \dot{u}_{\beta_1}(\underline{\theta}) - \frac{(\dot{u}_{\beta_1}(\underline{\theta}))^2}{2} - u_{\beta_1}(\underline{\theta})}{\alpha + \beta_1 \underline{\theta}} \right) \\ &= \underline{\theta}^{-1} \left[\left(\frac{\beta_2 \underline{\theta}}{\alpha + \beta_2 \underline{\theta}} \right) - \left(\frac{\beta_1 \underline{\theta}}{\alpha + \beta_1 \underline{\theta}} \right) \right] (2 - \ddot{u}_{\beta_2}(\underline{\theta})) > 0. \end{aligned}$$

Hence if $\ddot{u}_{\beta_1}(\underline{\theta}) \leq \ddot{u}_{\beta_2}(\underline{\theta})$ then there exists $\theta_1 > \underline{\theta}$ such that $\ddot{u}_{\beta_1}(\theta) < \ddot{u}_{\beta_2}(\theta)$ and $\dot{u}_{\beta_1}(\theta) < \dot{u}_{\beta_2}(\theta)$ for all $\theta \in (\underline{\theta}, \theta_1)$. Clearly there should be $\theta_2 \in (\theta_1, \bar{\theta})$ such that $\ddot{u}_{\beta_1}(\theta_2) = \ddot{u}_{\beta_2}(\theta_2)$, otherwise we would have $\dot{u}_{\beta_1}(\bar{\theta}) < \dot{u}_{\beta_2}(\bar{\theta})$, which is a contradiction. Therefore let $\theta^* \in (\theta_1, \bar{\theta})$ be the smallest element of this set such that $\ddot{u}_{\beta_1}(\theta^*) = \ddot{u}_{\beta_2}(\theta^*)$. Notice that we must have $\frac{d\dot{u}_{\beta_1}(\theta)}{d\theta} \Big|_{\theta^*} \geq \frac{d\dot{u}_{\beta_2}(\theta)}{d\theta} \Big|_{\theta^*}$ or equivalently

$$\begin{aligned} & \frac{d}{d\theta} \left[\left(\frac{1}{\alpha + \beta_1 \theta^*} \right) \left(\theta^* \dot{u}_{\beta_1}(\theta^*) - \frac{(\dot{u}_{\beta_1}(\theta^*))^2}{2} - u_{\beta_1}(\theta^*) \right) \right] \\ & \leq \frac{d}{d\theta} \left[\left(\frac{1}{\alpha + \beta_2 \theta^*} \right) \left(\theta^* \dot{u}_{\beta_2}(\theta^*) - \frac{(\dot{u}_{\beta_2}(\theta^*))^2}{2} - u_{\beta_2}(\theta^*) \right) \right], \end{aligned}$$

which holds if and only if

$$\begin{aligned} & \left(\frac{1}{\alpha + \beta_1 \theta^*} \right) \left[\theta^* - \dot{u}_{\beta_1}(\theta^*) \right] \ddot{u}_{\beta_1}(\theta^*) - \left(\frac{\beta_1}{(\alpha + \beta_1 \theta^*)^2} \right) \left(\theta^* \dot{u}_{\beta_1}(\theta^*) - \frac{(\dot{u}_{\beta_1}(\theta^*))^2}{2} - u_{\beta_1}(\theta^*) \right) \\ & \leq \left(\frac{1}{\alpha + \beta_2 \theta^*} \right) \left[\theta^* - \dot{u}_{\beta_2}(\theta^*) \right] \ddot{u}_{\beta_2}(\theta^*) - \left(\frac{\beta_2}{(\alpha + \beta_2 \theta^*)^2} \right) \left(\theta^* \dot{u}_{\beta_2}(\theta^*) - \frac{(\dot{u}_{\beta_2}(\theta^*))^2}{2} - u_{\beta_2}(\theta^*) \right). \end{aligned}$$

But notice that $\dot{u}_{\beta_1}(\theta^*) < \dot{u}_{\beta_2}(\theta^*) < \theta^*$ implies $\theta^* - \dot{u}_{\beta_1}(\theta^*) > \theta^* - \dot{u}_{\beta_2}(\theta^*) > 0$. This, $\ddot{u}_{\beta_1}(\theta^*) = \ddot{u}_{\beta_2}(\theta^*) > 0$ and $\left(\frac{1}{\alpha + \beta_1 \theta^*} \right) > \left(\frac{1}{\alpha + \beta_2 \theta^*} \right)$ imply

$$\left(\frac{1}{\alpha + \beta_1 \theta^*} \right) \left[\theta^* - \dot{u}_{\beta_1}(\theta^*) \right] \ddot{u}_{\beta_1}(\theta^*) > \left(\frac{1}{\alpha + \beta_2 \theta^*} \right) \left[\theta^* - \dot{u}_{\beta_2}(\theta^*) \right] \ddot{u}_{\beta_2}(\theta^*).$$

On the other hand, using $\left(\frac{\theta^* \dot{u}_{\beta_i}(\theta^*) - \frac{(\dot{u}_{\beta_i}(\theta^*))^2}{2} - u_{\beta_i}(\theta^*)}{\alpha + \beta_i \theta^*} \right) = 2 - \ddot{u}_{\beta_i}(\theta^*) > 0$, we immediately get

$$\begin{aligned} & - \left(\frac{\beta_1}{(\alpha + \beta_1 \theta^*)^2} \right) \left(\theta^* \dot{u}_{\beta_1}(\theta^*) - \frac{(\dot{u}_{\beta_1}(\theta^*))^2}{2} - u_{\beta_1}(\theta^*) \right) + \left(\frac{\beta_2}{(\alpha + \beta_2 \theta^*)^2} \right) \left(\theta^* \dot{u}_{\beta_2}(\theta^*) - \frac{(\dot{u}_{\beta_2}(\theta^*))^2}{2} - u_{\beta_2}(\theta^*) \right) \\ & = \frac{1}{\theta^*} \left[\left(\frac{\beta_2 \theta^*}{\alpha + \beta_2 \theta^*} \right) - \left(\frac{\beta_1 \theta^*}{\alpha + \beta_1 \theta^*} \right) \right] (2 - \ddot{u}_{\beta_1}(\theta^*)) > 0. \end{aligned}$$

Putting these together one gets $\frac{d\dot{u}_{\beta_1}(\theta)}{d\theta} \big|_{\theta^*} < \frac{d\dot{u}_{\beta_2}(\theta)}{d\theta} \big|_{\theta^*}$, a contradiction. Therefore we conclude that $\ddot{u}_{\beta_1}(\underline{\theta}) > \ddot{u}_{\beta_2}(\underline{\theta})$. This implies that there exists $\varepsilon > 0$ such that $\dot{u}_{\beta_1}(\theta) > \dot{u}_{\beta_2}(\theta)$ for all $\theta \in (\underline{\theta}, \underline{\theta} + \varepsilon)$. Assume towards a contradiction that $\dot{u}_{\beta_1}(\theta_*) = \dot{u}_{\beta_2}(\theta_*)$ for some $\theta_* < \bar{\theta}$ and let θ_* be the smallest element greater than $\underline{\theta} + \varepsilon$ satisfying this equality. We must have $\ddot{u}_{\beta_1}(\theta_*) \leq \ddot{u}_{\beta_2}(\theta_*)$. This and $\dot{u}_{\beta_1}(\theta_*) = \dot{u}_{\beta_2}(\theta_*)$ allows us to apply the first part of the proof and obtain a contradiction.

Case 2: $\beta < 0$. The proof is analogous to the case above and is omitted by brevity.

An increase in α decreases $u(\theta)$ for every type.

Keep β constant and take two solutions $u_{\alpha_1}(\theta)$ and $u_{\alpha_2}(\theta)$ for $\alpha_2 > \alpha_1$. First we show that $u_{\alpha_1}(\underline{\theta}) > u_{\alpha_2}(\underline{\theta})$. Assume towards a contradiction that $u_{\alpha_1}(\underline{\theta}) \leq u_{\alpha_2}(\underline{\theta})$ and notice that this implies $\ddot{u}_{\alpha_1}(\underline{\theta}) < \ddot{u}_{\alpha_2}(\underline{\theta})$ and hence there is $\theta_1 > \underline{\theta}$ such that for all $\theta \in (\underline{\theta}, \theta_1)$ we have $u_{\alpha_1}(\underline{\theta}) < u_{\alpha_2}(\underline{\theta})$ and $\dot{u}_{\alpha_1}(\theta) < \dot{u}_{\alpha_2}(\theta)$. Let θ_* be the smallest element θ of $[\theta_1, \bar{\theta}]$ such that $\dot{u}_{\alpha_1}(\theta) = \dot{u}_{\alpha_2}(\theta)$. We must have $\ddot{u}_{\alpha_1}(\theta) \geq \ddot{u}_{\alpha_2}(\theta)$, but then

$$\begin{aligned} \ddot{u}_{\alpha_2}(\theta_*) & = 2 - \left(\frac{1}{\alpha_2 + \beta \theta^*} \right) \left(\theta^* \dot{u}_{\alpha_2}(\theta^*) - \frac{(\dot{u}_{\alpha_2}(\theta^*))^2}{2} - u_{\alpha_2}(\theta^*) \right) \\ & = 2 - \left(\frac{1}{\alpha_2 + \beta \theta^*} \right) \left(\theta^* \dot{u}_{\alpha_1}(\theta^*) - \frac{(\dot{u}_{\alpha_1}(\theta^*))^2}{2} - u_{\alpha_1}(\theta^*) \right) + \left(\frac{1}{\alpha_2 + \beta \theta^*} \right) (u_{\alpha_2}(\theta^*) - u_{\alpha_1}(\theta^*)) \\ & > 2 - \left(\frac{1}{\alpha_1 + \beta \theta^*} \right) \left(\theta^* \dot{u}_{\alpha_1}(\theta^*) - \frac{(\dot{u}_{\alpha_1}(\theta^*))^2}{2} - u_{\alpha_1}(\theta^*) \right) + \left(\frac{1}{\alpha_2 + \beta \theta^*} \right) (u_{\alpha_2}(\theta^*) - u_{\alpha_1}(\theta^*)) > \ddot{u}_{\alpha_1}(\theta_*), \end{aligned}$$

a contradiction.

Next, suppose towards a contradiction that there exists θ such that $u_{\alpha_2}(\theta) - u_{\alpha_1}(\theta) > 0$, and then take $\theta^* \in \operatorname{argmax}_{\theta} u_{\alpha_2}(\theta) - u_{\alpha_1}(\theta)$. Notice that whether $\theta^* < \bar{\theta}$ or $\theta^* = \bar{\theta}$, we must have $u_{\alpha_2}(\theta^*) - u_{\alpha_1}(\theta^*) > 0$ and $\ddot{u}_{\alpha_2}(\theta^*) - \ddot{u}_{\alpha_1}(\theta^*) \leq 0$, but then the same argument as above implies $\ddot{u}_{\alpha_2}(\theta_*) > \ddot{u}_{\alpha_1}(\theta_*)$, which leads to a contradiction. \blacksquare

Proposition 7. *There exists at most one pure-strategy equilibrium.*

Proof of Proposition 7. Follows from Facts 1 and 2 below. \blacksquare

Fact 1. *If there exists an equilibrium (u_l^1, u_h^1) in which no IC constraint binds then this is the unique pure strategy equilibrium.*

Proof of Fact 1. Using the strict logconcavity of the profit function, it is trivial to see that there is at most one equilibrium in which no IC constraint binds and in which case this is in pure strategies. Hence we show that there is no other equilibrium in which one IC binds. Recall that $\Delta(u_l, u_h) := u_h - u_l$. Notice that this equilibrium satisfies $u_h^1 > 0$, and

$$S_l^e - u_l^1 - \lambda t_l \leq 0, \quad (= 0 \text{ if } u_l^1 > 0) \quad (24)$$

and

$$S_h^e - u_h^1 - \lambda t_h = 0. \quad (25)$$

First suppose that there exists an equilibrium (u_l^2, u_h^2) in which IC_l binds:

$$S_l^e - u_l^2 + \lambda t_l \left(\frac{p_h}{p_l} \frac{\partial S_h}{\partial u_l} (\Delta(u_l^2, u_h^2)) - 1 \right) \leq 0, \quad (= 0 \text{ if } u_l^2 > 0) \quad (26)$$

and

$$S_h(u_l^*, u_h^*) - u_h^2 + \lambda t_h \left(\frac{\partial S_h}{\partial u_h} (\Delta(u_l^2, u_h^2)) - 1 \right) = 0. \quad (27)$$

Notice that (25) and (27) imply $u_h^2 < u_h^1$, while (24) and (26) imply $u_l^2 \geq u_l^1$ implying $\Delta(u_l^2, u_h^2) < \Delta(u_l^1, u_h^1)$, a contradiction.

Next suppose that there exists an equilibrium (u_l^3, u_h^3) in which IC_h binds:

$$S_l(u_l^3, u_h^3) - u_l^3 + \lambda t_l \left(\frac{\partial S_l}{\partial u_l} (u_l^3, u_h^3) - 1 \right) \leq 0, \quad (= 0 \text{ if } u_l^3 > 0) \quad (28)$$

and

$$S_h^e - u_h^3 + \lambda t_h \left(\frac{p_l}{p_h} \frac{\partial S_l}{\partial u_h} (u_l^3, u_h^3) - 1 \right) = 0. \quad (= 0 \text{ if } u_h^3 > 0) \quad (29)$$

Notice that (29) and (25) imply $u_h^3 > u_h^1$, while (28) and (26) imply $u_l^3 \leq u_l^1$ delivering $\Delta(u_l^3, u_h^3) > \Delta(u_l^1, u_h^1)$, which is a contradiction. ■

Fact 2. *If there exists an equilibrium in which the IC_l constraint binds then this is the unique pure-strategy equilibrium.*

Proof of Fact 2. Let (u_l^2, u_h^2) an equilibrium in which IC_l binds and (u_l^3, u_h^3) an equilibrium in which IC_h binds. Using (27) and (29) we $u_h^3 > u_h^2 > 0$. On the other hand (28) and (26) imply $u_l^3 \leq u_l^2$. Therefore we have $\Delta(u_l^3, u_h^3) > \Delta(u_l^2, u_h^2)$, which is a contradiction. ■