COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Post Office Box 208281
New Haven, Connecticut 06520–8281

COWLES FOUNDATION DISCUSSION PAPER NO. 1199

Price Competition for an Informed Buyer

by

Giuseppe Moscarini

and

Marco Ottaviani

October 1998

Price Competition for an Informed Buyer*†

Giuseppe Moscarini‡ **Marco Ottaviani**§

September 1998

Abstract

We investigate the outcomes of simultaneous price competition in the presence of private information on the demand side. Each of two sellers offers a different variety of a good to a buyer endowed with a private binary signal on their relative quality. We analyze how the unique equilibrium of the game changes as a function of the (common) prior belief on the relative quality of the goods and the precision of the private information of the buyer. Competition is fierce, and the buyer enjoys high rents, when the prior belief is biased in favor of one good and private signals are not very informative: the ex ante superior seller cannot resist the temptation to clear the market, and triggers an aggressive response by the competitor. When instead the distribution of ex post valuations is highly spread, due to a vague prior belief and strong signals, the sellers become local monopolists and extract high rents from the buyer. We provide a full characterization of the mixed-strategy equilibrium which arises when the two goods are mildly differentiated ex post. Overall, the market-clearing temptation effect destroys the monotonicity and convexity of the equilibrium profit of a seller in the prior belief. As a consequence, a competing seller does not necessarily benefit from revelation of public information, sometimes even if biased in its favor.

---

* We thank: Abhijit Banerjee, Glenn Ellison, and Bengt Holmström for helpful advice; Phil Reny and Leo Simon for insightful correspondence and conversation; Ken Binmore, Tilman Börgers, Jacques Crémer, Ian Jewitt, Alessandro Lizzeri, Fausto Panunzi, Patrick Rey, Larry Samuelson, Avner Shaked, Peter Sørensen, Jeroen Swinkels, Juuso Valimaki, and the audience at IGIER Colloquia, Warwick Conference on Dynamic Games, Fourgeaud Seminar in Paris, Bristol, and Tilburg for fruitful comments. Part of this paper is contained in Chapter 2 of Ottaviani's MIT doctoral dissertation.

† *JEL Classification*: D43 (Oligopoly), D44 (Auctions), D82 (Asymmetric and Private Information), L15 (Information and Product Quality).

‡Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven CT 06511, USA. Tel: +1-203-432-3596. Fax: +1-203-432-5779. E-mail: `gm76@pantheon.yale.edu`.

§Department of Economics and ELSE, University College London, Gower St., London WC1E 6BT, UK. Tel: +44-171-504-5242. Fax: +44-171-916-2775. E-mail: `m.ottaviani@ucl.ac.uk`.

1. Introduction

Martha is determined to hire the best decorator to revamp her vast town house in Central London with stucco curved ceilings, stone-finished walls and mosaic floors. The classical design of Sanderson — a decorator with a long-standing tradition in London — has clearly an edge on the funkier design of Conran. Each decorator privately gives her a detailed description and demonstration of the services to be provided and quotes an inclusive price. Martha then decides which decorator to hire, if either, on the basis of this information. In this and other similar circumstances the price-taking party knows more accurately than the competing price-setters the match of her preferences (or technology) with the intrinsic characteristics of the good or service exchanged. As further examples, the user of an input of production is likely to be better informed than the suppliers on its quality, relative to the available alternatives, and a sub-contractor knows better the cost of the specific product or service required by competing contractors. In the case of sellers competing for an informed buyer, what is the outcome of price competition among the sellers? How much information should the buyer acquire? To answer these questions we extend the classic Bertrand (1883) model of price competition to asymmetric settings with private information on the demand side.

The problem of static competition for a privately informed buyer is a natural building block for dynamic models of strategic pricing with learning in the presence of private information. While Bergemann and Valimaki (1996), (1997), and Felli and Harris (1996) analyze the dynamics of price competition as public information spreads, the final aim of our research project is to study the case with private information. In such dynamic models the level of differentiation among the goods is the state variable which evolves with the accumulation of information on relative quality.

In our static model ex-post differentiation between the sellers is the result of prior vertical differentiation and private information of the buyer. First, sellers may differ *ex ante*, when it is common knowledge in the market that one variety of the good (Sanderson) is more likely to be better than the other (Conran) before the customer receives any private information on their relative quality. Second, sellers appeal *ex-post* in different ways to the buyer, depending on the private signal (description and demonstration of the services to be provided) observed before the purchase decision is made. In this paper we focus mostly on the prototypical case of *two* sellers competing for a single buyer, who observes a private *binary* signal on the *relative* quality of the two goods offered. Equilibrium prices and profits of the sellers are characterized as functions of two crucial parameters of the model: the common prior belief on the relative quality of the two goods, and the precision of the buyer's private information. The equilibrium is unique, so that we may illustrate unambiguous comparative statics predictions on equilibrium strategies and payoffs for changes in the two parameters.

The equilibrium strikes a balance between two tensions: each duopolist on the one hand wishes to extract the maximum rent from the buyer for given probability of selling, while on the other is tempted to reduce the price to increase the probability of selling at the expense of the competitor. Slightly undercutting the price of the competitor can yield a discrete gain in quantity demanded and therefore in expected profits, so that the payoff functions of the sellers are discontinuous. Given the close connection of the games, the equilibrium constructed shares many common features with that found by Osborne and Pitchik (1987) in the Hotelling (1929) model with uniform distribution of consumers. In two extreme regions of parameters the equilibrium is in pure strategies, while the sellers play a mixed-strategy in the remaining intermediate region.

Firstly, a unique pure-strategy equilibrium arises for a sufficiently balanced prior belief on the relative quality of the two goods and a sufficiently high precision of the private signal of the customer. In this case, final differentiation between the two varieties is strong and the customer becomes very keen to buy one of them depending on the realization of the private signal received. In turn, each supplier only targets the customer with favorable information for her own good, since it would be too costly to steal from the competitor the customer who received unfavorable information. This *separating equilibrium* is characterized by little competitive pressure on prices and high profits for both sellers, who become local monopolists and leave no rents to the buyer.

Secondly, when the prior quality perceived by the market is biased enough for one seller and the precision of the private signal is low enough, the weak power of private information cannot reverse the strong ex-ante inclination of the customer for one of the two goods. Competition results then in a *pooling equilibrium*; the seller favored by the prior belief covers the entire market by posting a limit price which excludes the competitor. This price is necessarily low, relative to the level of perceived quality, because it must induce the customer to buy even after receiving an unfavorable signal. Competition has therefore a tendency to become fiercer and to lead to lower prices as uncertainty on the relative quality decreases. In this region of parameters both types of customer enjoy rents.

Finally, in cases of mild final differentiation, both firms randomize over the continuous price space in the unique equilibrium, in *mixed strategies* by the sellers. Moving from the separating to the pooling region the equilibrium randomizations put more weight on low prices, as the seller favored ex ante has more incentive to undercut the competitor. Our constructive characterization of the mixed strategy equilibrium can be useful in other games with a continuum of actions and discrete private information.

The *comparative statics* of equilibrium with respect to changes in the two parameters reveals a natural continuity of equilibrium payoffs across different equilibrium regimes. More importantly, it offers interesting insights on the value of information in a strategic setting with endogenous pricing. First, consider changes in the prior belief, achieved by revelation of public information to both buyer and seller. Such information allows for better decisions and is therefore socially beneficial, if one abstracts from distributional issues by keeping prices fixed. The best-known result in a strategic setting is due to Milgrom and Weber (1982): in a second-price auction with affiliated valuations of the buyers revealing public information is always beneficial to the seller. A simple implication of this result is that the profit function of a monopolist who sells to a privately informed buyer is convex in the common prior belief. Our main finding is that this general convexity property is lost whenever another differentiated seller who competes in price is introduced. In particular, a seller dislikes revelation of public information when the prior belief is intermediate and private information of the buyer is strong. The reason is again that extreme beliefs lead the sellers to compete more fiercely. This non-convexity contrasts not only with the results for the case of non-strategic pricing, but also with those obtained for learning without private information. We also find that the payoff of the ex-ante superior seller is not monotonic in the prior belief, in contrast to the monopoly case. Finally, the sum of the payoffs of the sellers is non-monotonic in the level of ex-ante differentiation, in contrast to the case without private information.

Second, consider changes in the precision of the private information of the buyer. In contrast to the prediction in a fixed-price environment, the payoff of the buyer is non-monotonic in the precision of her own private signal. The buyer is strictly worse off by acquiring private information beyond a certain level, because the sellers can extract more rents in equilibrium the greater the ex-post differentiation of valuations. Similarly, the payoff of a seller is non-monotonic in the precision of the buyer's private information, even though more precise signals raise the total surplus to be shared among the society of sellers and buyer.

When the sellers compete ex ante, the buyer does not necessarily purchase the ex-post superior good, as would be required to achieve allocative efficiency. In particular, for mild final differentiation (in the mixed strategy region and in part of the separating region) the equilibrium is inefficient because the ex-ante superior seller has an incentive to keep the price high in order to extract more rent from the consumer.

The paper proceeds as follows: Section 2 briefly surveys the literature. Section 3 introduces our basic duopoly model of competition for an informed buyer and compares it to that of Hotelling (1929). Section 4 characterizes the equilibrium for all parameter values. Our main results on the value of information for the buyer and the sellers are derived and discussed in Section 5. Section 6 deals with the robustness of our analysis and results to changes in the assumptions. Section 7 concludes.

2. Relation to the Literature

Extreme predictions are obtained very easily in the Bertrand model with identical sellers, but departing from the classical case can lead to formidable technical complications. This could explain why this problem is relatively unexplored. Because private information on the demand side gives rise to ex-post differentiation between the two sellers, our model has many formal analogies (and similar technical problems) to a Hotelling (1929) pricing game for given locations of the sellers. In our setting the prior belief and the precision of the signal affect the distribution of the buyer's valuation: a prior more biased in favor of one seller shifts the distribution closer to that seller, and more accurate information of the buyer corresponds to a mean preserving spread in the distribution. In a similar vein, Gabszewicz and Grilo (1992) study price competition conditional on quality in a duopoly market where firms sell vertically differentiated products and consumers have heterogeneous beliefs on quality. In order to tackle the problem of information acquisition, we instead perform unconditional analysis in a market for a single buyer with belief derived from a common prior. The restrictions on the corresponding reduced-form Hotelling model provided by our parametrization are discussed throughout the paper.

Broecker (1990) analyzes the (mixed strategy) equilibrium in a competitive market where each bank performs a binary test on applicant firms and offer credit conditional on the realization of this test. In his setting each competitor has an independent (private) signal on the firm, while we consider the specular case where the price-taking agent alone has private information. Mixed strategy equilibria have also been constructed in rather different models of price competition. For example, Rosenthal and Weiss (1983) characterize the mixed-strategy equilibrium in Spence's model of competition in the presence of signalling. Their equilibrium construction extends to other models of competition with discrete asymmetric information, like the Rothschild and Stiglitz model of competition in the insurance market. The features of their equilibrium are completely different from ours, because in those models the goods offered by different competitors are identical, while in our setting they are differently appealing to different types of consumers. As a further example, Varian (1980) models sellers who engage in sales behavior in an attempt to price discriminate between informed and uninformed consumers in a Butters (1977)-like world.

Competing mechanisms for selling to informed buyers have received recent attention in auction theory and industrial organization. For instance, McAfee (1993) constructs a dynamic model of price formation where sellers compete in designing mechanisms to sell identical objects to buyers. In our setting sellers offer instead differentiated goods and the buyer has private information on their relative quality. Stole (1993) studies price discrimination with non-linear pricing in differentiated oligopoly with private information of consumers on preference for brand and quality. As discussed in Section 6.1, in our model the linearity of the preferences of the buyer do not allow any scope for action on the quantity dimension even when sellers are allowed to offer menus of lotteries.

Our model of price competition can also be interpreted as a first-price procurement auction where the buyer decides which good to buy depending on the price bids as well as the realization of a private signal on the relative quality of the goods offered. Manelli and Vincent (1995) study optimal procurement mechanisms in environments where the valuation of the buyer depends on the sellers' private information. In their setting a procurement auction may be suboptimal for the buyer, being dominated by a credible take-it-or-leave-it offer to a seller selected at random. Similarly in our setting, where only the buyer has private information, a credible offer of a price equal to zero to the ex-post superior-quality seller would maximize the payoff of the buyer and achieve the efficient allocation.[1] Though our first-price auction is not the optimal procurement mechanism for the buyer in this environment, it arises naturally when the bargaining power does not rest with the informed party. In Section 6.1 we show that in our environment with linear preferences simple price competition results when competing sellers optimally design the mechanism and are allowed to offer menus of lotteries, provided that they cannot condition their offers to those of the competitor.

Finally, our work relates to the literature on experimentation in oligopoly. For instance, Harrington (1995) extends the model of learning of the market demand by a monopolist to an oligopoly setting. Closer in spirit to us, Schlee (1996) considers the value of public information in a model where buyer and seller share the same belief on quality. In his setting the buyer demands multiple units and does not possess any private information.

3. Model

Supply. Consider the following one-shot model of simultaneous price competition for an informed buyer. On the supply side of the market there are two sellers competing simultaneously on prices. Each duopolist, denoted by $j \in \{0, 1\}$, posts price P_j for her variety of the good and commits to sell at that price if the buyer agrees. The marginal cost of both suppliers is set equal to zero for convenience of notation. Each firm is risk-neutral and maximizes expected profits.

Demand. On the demand side there is a single risk-neutral buyer with a unit demand for an indivisible good. The payoff in case of no purchase (action $a_\emptyset$) is 0. There are two states of nature, ω_0 and ω_1, indicating the superior good. Since good 1 is better than good 0 in state ω_1, while the opposite is true in state ω_0, we assume for convenience that

[1] This result holds since there is no private information of the sellers and no entry costs to become a seller.

the (gross of price) payoff of purchasing good i (action a_i) in state ω_j is 1 if $i = j$ and 0 if $i \neq j$. Relative quality indicates the match of the preferences of the buyer with the intrinsic characteristics of the good. The buyer maximizes the expected valuation net of the price paid.

Information. The state ω is unknown to the sellers and the buyer who share the same (common knowledge) prior belief $q = \Pr(\omega_1)$. Prior to purchase, the buyer observes the realization of a private signal of bounded precision on the relative quality of the two varieties and the prices posted by both sellers. In particular, we consider the case where the signal $\sigma \in \{\sigma_0, \sigma_1\}$ is binary with conditional probability distributions:

$$\Pr(\sigma_i|\omega_j) = \begin{cases} \alpha & \text{for } j = i \\ 1-\alpha & \text{for } j \neq i \end{cases} \tag{3.1}$$

with $i, j \in \{0, 1\}$. Notice that for simplicity we are considering the symmetric case where $\Pr(\sigma_i|\omega_i) = \alpha$ for $i = 0, 1$ is the probability that the buyer receives the "correct" signal. Without loss of generality we restrict attention to $\alpha \in [1/2, 1]$, since $\alpha \in [0, 1/2]$ would be equivalent to relabelling the signals. Since decreasing α corresponds to garbling this binary signal structure, α parametrizes the precision (or "quality") of the buyer's private information. The quality of the private signal is *bounded* whenever $\alpha < 1$ and the signal is informative for $\alpha \neq 1/2$. The private signal received by the customer is not observed by the sellers, but the conditional distributions of the signal in (3.1) are common knowledge.[2]

Bayesian Updating. Let $f_i(q, \alpha) \equiv \Pr(\omega_1|\sigma_i)$ be the buyer's private posterior belief that the state is ω_1 after observing the signal realization σ_i. Bayes' rule yields

$$f_0(q,\alpha) = \frac{(1-\alpha)q}{\alpha(1-q) + (1-\alpha)q}, \quad f_1(q,\alpha) = \frac{\alpha q}{\alpha q + (1-\alpha)(1-q)}. \tag{3.2}$$

The updated belief represents, given the zero/one payoffs, the customer's expected valuation for good 1 in monetary terms. The ex-ante valuations are q for good 1 and $1 - q$ for good 0, and the ex-post valuations of the buyer with signal σ_i are f_i for good 1 and $1 - f_i$ for good 0. Notice that $\partial f_0/\partial q > 0$, $\partial f_1/\partial q > 0$, $\partial^2 f_1/\partial q^2 < 0 < \partial^2 f_0/\partial q^2$, and $\partial f_0/\partial \alpha < 0 < \partial f_1/\partial \alpha$.

Game. Figure 1 illustrates the extensive form of the game. First, both firms quote simultaneously prices to which they precommit to sell to the customer. Then nature

[2]The binary signal formulation is widely adopted, being the simplest one entailing information of bounded precision (see e.g. Broecker's (1990) study of competition with adverse selection). The quality of our analysis should not change with alternative formulations with finite number of signals. We refer to Section 6.2 for a discussion of alternative formulations of the model with a continuous signal structure.

determines the realization of the signal according to (3.1). The true state of nature is not known to the firms, so that their common probability assessment that signal σ_i is received by the customer is

$$\Pr(\sigma_i|q,\alpha) \equiv q\Pr(\sigma_i|\omega_1) + (1-q)\Pr(\sigma_i|\omega_0)\,. \tag{3.3}$$

The customer observes the signal realization σ_i and updates the belief from q to f_i. In the language of Bayesian games, we will refer to the customer who has received signal σ_i as to *type-i* customer, since the realization of the signal is not known to the sellers. The customer then compares her updated valuation with the prices quoted and decides whether to buy and from which firm to buy. The payoff to a firm when selling is equal to the price charged and the payoff of the customer is equal to the valuation for the good bought minus the price paid for it. The order in the payoff vector in the figure is: seller 0, seller 1, and the buyer.

Buyer's Behavior. The customer chooses the good yielding the highest expected payoff net of the price (*incentive compatibility* constraint), provided that the net payoff is non negative because of the outside option of not purchasing (*individual rationality* constraint). The net valuations of type-i buyer are $f_i - P_1$ for good 1 and $1 - f_i - P_0$ for good 0, so that good 1 is preferred to good 0 when

$$f_i - P_1 > 1 - f_i - P_0, \tag{$IC_{i,1}$}$$

where $IC_{i,j}$ stands for incentive compatibility for buyer of type i to buy good j. By reversing the inequality one obtains the $IC_{i,0}$ constraint. Type-i customer is exactly indifferent for prices

$$P_1 = 2f_i - 1 + P_0 \tag{IC_i}$$

at which IC_i binds. With some abuse of notation, $P_j \equiv IC_i(P_{1-j})$ stands conveniently for the price of firm j corresponding to P_{1-j} through the IC_i constraint, precisely $IC_i(P_0) = 2f_i - 1 + P_0$ and $IC_i(P_1) = 1 - 2f_i + P_1$. The maximum willingness to pay for good j by type$-i$ customer is determined by the individual rationality constraint $IR_{i,j}$

$$f_i - P_1 \;\geq\; 0, \tag{$IR_{i,1}$}$$

$$1 - f_i - P_0 \;\geq\; 0. \tag{$IR_{i,0}$}$$

In summary, type-i customer buys good 1 if $P_1 < \min\langle 2f_i - 1 + P_0, f_i\rangle$, and good 0 if $P_0 < \min\langle 1 - 2f_i + P_1, 1 - f_i\rangle$. Figure 2a represents the IC's and IR's constraints in the P_0, P_1 space. The $IR_{i,1}$ lines are horizontal and $IR_{i,0}$ vertical. The IC_i lines have both unit slope. As illustrated in the figure, $IC_{1,1}$ is binding for $P_0 < 1 - f_1$ and $IR_{1,1}$ is binding otherwise. $IC_{1,0}$, rather than $IR_{1,0}$, is instead always binding in the relevant range of prices.

Differentiation. To understand the role of the parameters q and α, it is useful to consider the level of *ex-ante differentiation* (or asymmetry) $|q-1/2|$ as a measure of *vertical differentiation* and $\Delta f \equiv f_1 - f_0 > 0$ — the difference of valuations for good 1 between the two types of consumers — as a measure of *horizontal differentiation*. Finally, *ex-post differentiation* is measured by $\sum_{i=0}^{1} \Pr(\sigma_i)\,|f_i - 1/2|$. As q increases above $1/2$, the valuations f_0 and f_1 for good 1 of both consumer's types increase, while the spread in the distribution of valuations Δf decreases. The reduction of the degree of horizontal differentiation is a by-product of Bayesian updating; a signal of given precision is less informative the more concentrated the prior probability distribution, here the more extreme q is. As of the distribution of demand at the two locations, q raises $\Pr(\sigma_1)$ and reduces $\Pr(\sigma_0)$. An increase in α spreads the distribution of the posterior valuations further apart by increasing f_1 and reducing f_0, thereby augmenting Δf, while still raising $\Pr(\sigma_1)$ and reducing $\Pr(\sigma_0)$ for $q > 1/2$ (the opposite for $q < 1/2$). Overall, an increase in α entails a mean-preserving spread of the distribution of posterior valuation. Finally, ex-post differentiation is large either for high ex-ante differentiation $|q-1/2|$ even when the spread α is small, or for low ex-ante differentiation but large spread.

Comparison to Hotelling. Figure 2b gives a reduced-form representation of this model in the Hotelling line. Seller 0 is located at the origin and seller 1 at the other end of the segment of unit length. With probability $\Pr(\sigma_1)$ the consumer (of type 1) is located at f_1, and with complementary probability is of type 0 at f_0. Probability can also be interpreted as mass of demand at these locations. As usual in Hotelling models of product differentiation, the utility of consuming good j for the consumer of type i is decreasing in the distance $d(i,j)$ of that consumer from seller j. In particular, in our model the utility of the consumer of type i is equal to $1-d(i,j)-P_j$ when buying from seller j at price P_j, and equal to 0 when not buying. The structure of our model imposes restrictions on how to perform comparative statics with respect to the probability distribution of the valuation of the buyer. These restrictions would not have been obtained if one departed directly from a reduced-form Hotelling model.

4. Equilibrium

Each duopolist wishes to extract the maximum rent from the buyer for given expected quantity demanded and at the same time is tempted to steal demand from the competitor. The *Bayes-Nash equilibrium* of the game strikes a balance between these two forces, depending on the parameters of the model: the quality of private information $\alpha \in [1/2, 1]$ and the prior belief $q \in [1/2, 1]$ represented in Figure 3.[3] Notice that the game is symmet-

[3]The belief q is the natural state variable in dynamic extensions of the model discussed in Section 7.

ric with respect to $q = 1/2$, so that we need only consider $q \geqslant 1/2$. Once the equilibrium outcome is characterized for a fixed parameter configuration (q, α), we analyze how it changes with the parameters in Section 5.

Before proceeding to the characterization, we briefly discuss the issue of equilibrium existence which is a non-trivial problem in such continuous games with discontinuous payoffs. In this game the payoff function of each seller jumps at prices of indifference for the buyer, i.e. along the IC constraints where *ties* arise for each buyer type. As a consequence, the sum of payoffs is not upper semi-continuous (USC), violating the second Dasgupta and Maskin (1986) sufficient condition for existence. For USC we would need the indifferent buyer to always break the tie in favor of the high-price firm.[4] We are able to avoid such technical difficulties by constructing the equilibrium for each game corresponding to any parameter configuration (q, α). In our model ties happen with probability zero in equilibrium, except for parameter configurations where a pure strategy equilibrium is immediately found. This irrelevance of ties in equilibrium may generalize to richer games with continuum of actions and private information.

4.1. Best Replies

As a preliminary step toward the characterization of the equilibrium, we construct the best reply functions of the sellers. Given a price P_0 posted by seller 0, only three strategies are not patently dominated for seller 1: not selling at all, selling only to type-1 customer at the *separating* price

$$P_1^S(P_0) \equiv \begin{cases} 2f_1 - 1 + P_0 & \text{for } P_0 \leq 1 - f_1 \\ f_1 & \text{for } P_0 \geq 1 - f_1, \end{cases} \tag{4.1}$$

or to both customer types at the *pooling* price

$$P_1^P(P_0) \equiv 2f_0 - 1 + P_0. \tag{4.2}$$

The best reply is the strategy which achieves $\max \langle 0, \Pr(\sigma_1) P_1^S(P_0), P_1^P(P_0) \rangle$, where the first option corresponds to not selling at all, the second to separating and the third to pooling. At the (unique) *switching price* $\hat{P}_j$ the best reply function of firm $1-j$ jumps down from the separating $P_{1-j}^S(\cdot)$ to the pooling price $P_{1-j}^P(\cdot)$, so that $\Pr(\sigma_{1-j}) P_{1-j}^S(\hat{P}_j) \equiv P_{1-j}^P(\hat{P}_j)$.

[4] This conclusion is not new in Hotelling games with continuous distributions of customers, but our introduction of discrete amounts of private information (or equivalently of a discrete distribution of demand) makes the identity of the firm to be favored by the indifferent customer depend on the specific tying price pair. This endogeneity of the tie-breaking rule is reminiscent of the general result of Simon and Zame (1990), though in our game the tie-breaking rule is truly part of the strategy of the buyer.

4.2. Equilibrium Characterization

Before proceeding to the general characterization of the equilibrium, we introduce four special cases of this game as simple reference points. First, $\alpha = 1$ corresponds to a buyer who is perfectly informed ex post, discussed at the end of Section 4.2.1. Second, when $q = 1$ the buyer is perfectly informed ex ante, a case we discuss at the end of Section 4.2.2. Third, when the sellers are ex-ante identical ($q = 1/2$), we have a perfectly symmetric version of our model, discussed in Section 4.2.5. Fourth, when there is no private information on the demand side ($\alpha = 1/2$), we are back to the Bertrand case with heterogeneous suppliers, an important benchmark for our analysis. In this case existence of equilibrium is guaranteed by requiring the indifferent customer to buy from the higher-quality seller. There is a continuum of Nash equilibria where firm 0 posts $P_0 \in [1-2q, 0]$ without being able to sell, and firm 1 sells at price $P_1 = 2q - 1 + P_0 \in [0, 1-2q]$ to the indifferent buyer. All these equilibria, other the one with $P_0 = 0$ and $P_1 = 2q - 1$, are usually disregarded as unreasonable because the low-quality seller plays weakly dominated strategies. In order to exclude the undesirable equilibria, Bergemann and Valimaki (1996) require equilibria to be *cautious*, in the sense that the non-selling firm would be indifferent between selling and not in equilibrium.[5] In the unique cautious equilibrium prices are $P_0(q, 1/2) = 0$, $P_1(q, 1/2) = 2q - 1$, the (indifferent) consumer buys from seller 1 with a net payoff of $V_B(q, 1/2) = 1 - q$, and profits of the sellers are $V_0(q, 1/2) = 0$, $V_1(q, 1/2) = 2q - 1$.

Our investigation now heads for the unexplored territory of asymmetric games with private information. In the parameter space (q, α) there are two regions (separating and pooling) with equilibrium in pure strategy, and one with a mixed strategy equilibrium.

4.2.1. Separating Equilibrium

When both switching prices exceed the maximum valuations for the two goods, the best reply correspondences of the two firms cross at the corner point $P_0 = 1 - f_0, P_1 = f_1$. The equilibrium is *separating* because the customer buys from a different seller depending on the realization of the signal. See Figure 4 for a representation in the price space. The no-deviation condition for firm 1 requires that pooling both types of buyer is less profitable than selling only to the ex-post favorable customer at the separating price, $2f_0 - 1 + P_0 = f_0 \leq \Pr(\sigma_1) f_1$, or equivalently $q < q^S(\alpha) \equiv (\alpha^2 + \alpha - 1)/[\alpha(2\alpha - 1)]$ with $dq^S/d\alpha > 0$. The no-deviation condition for seller 0, equivalent to $q \geq 1 - q^S$, is implied by $q \geq q^S$ for $q \geq 1/2$. The region of parameters where the unique equilibrium is separating corresponds to the area to the south-east of the lowest curve in Figure 3. For this equilibrium to exist, it is necessary that the quality of private information be large

[5] Uniqueness of equilibrium can be also obtained by eliminating strategies which are dominated according to the definition given by Börgers (1992) at page 168.

enough ($\alpha \geq 2/3$).

In the separating equilibrium region the sellers are weakly ex-ante differentiated, but strongly ex-post differentiated. Although both suppliers are ex ante in the race, the buyer has a strong ex post preference for the good favored by the signal realization. Each supplier anticipates this fact and targets only the customer with favorable information for its own good, resisting the temptation to steal from the competitor the customer with unfavorable information. In sharp contrast with the standard Bertrand paradox, in this equilibrium there is no competitive pressure on prices. The sellers become "local monopolists" and make high profits by fully extracting the customer's surplus.[6] In summary:

Proposition 1 (Separating equilibrium). *For $q < q^S(\alpha)$ the unique equilibrium is separating: the buyer who receives signal σ_i purchases from seller i, prices are $P_0(q,\alpha) = 1 - f_0(q,\alpha), P_1(q,\alpha) = f_1(q,\alpha)$, the sellers enjoy expected profits of $V_0(q,\alpha) = \alpha(1-q), V_1(q,\alpha) = \alpha q$ by extracting the entire rent of the buyer, who is left with expected payoff $V_B(q,\alpha) = 0$.*

In the special case when the buyer is perfectly informed ex post ($\alpha = 1$), $P_0(q,1) = P_1(q,1) = 1$, $V_0(q,1) = 1-q, V_1(q,1) = q$ and the buyer surplus is $V_B(q,1) = 0$. Compared to the case of no private information ($\alpha = 1/2$), it can already be seen that the buyer is worse off when perfectly informed, though the sum of the payoffs of buyer and sellers is highest at $\alpha = 1$. The increased level of competition among sellers destroys the buyer's incentives for acquisition of socially valuable information.

4.2.2. Pooling Equilibrium

When the switching price of one firm is below the marginal cost (e.g. $\hat{P}_0 \leq 0$), the only equilibrium is *pooling* on the good sold by the other firm, in this case good 1. In this equilibrium the customer buys from seller 1 regardless of the realization of the signal. Consider the prices $P_0 = 0, P_1 = IC_0(0) = 2f_0 - 1 > 0$, at which the indifferent type-0 buyer breaks the tie in favor of the high price seller 1. See Figure 5 for an illustration in the price space. Clearly, seller 0 has no profitable deviation since any non-negative price would not sell. The separating deviation for firm 1 is $IC_1(0) = 2f_1 - 1$, where the price couple $(0, IC_1(0))$ is along the IC_1 constraint and the $IR_{1,1}$ constraint is satisfied with strict inequality. The pooling deviation for seller 1 is not profitable provided that $IC_0(0) \geq \Pr(\sigma_1) IC_1(0)$, since the separating price $IC_1(0)$ sells with probability $\Pr(\sigma_1|q)$, while the pooling price $IC_0(0)$ sells with probability one. In the limit as q tends to 1,

[6] Full rent extraction only arises because the number of sellers (goods) is equal to that of buyer types (signal realizations). With non-binary signals, some buyer types would enjoy rents even in a separating equilibrium. These rents will be smaller the sharper ex post differentiation and type separation.

f_1 and f_0 both converge to 1, so that the separating price converges to the pooling one, while the probability of selling at the separating price tends to α. For q large enough it is then optimal for seller 1 to charge the pooling price, thereby selling with probability 1. A pooling equilibrium of this sort exists if only if $q \geq q^P(\alpha)$, where $q^P(\alpha)$ is the largest root of

$$2f_0(q,\alpha) - 1 = q - (1-\alpha), \tag{4.3}$$

with $q^P(\alpha) \in (\alpha, 1)$ and $dq^P/d\alpha > 0$. The pooling equilibrium region corresponds to the area to the north-west of the highest curve in Figure 3. Similarly to what we have done in the case without private information at the beginning of Section 4.2, we focus on the unique cautious equilibrium.[7]

In the pooling region the prior quality q perceived by the market is biased enough for one seller and the private signal precision α is low enough, that the final result of the strong ex-ante heterogeneity and a mild ex-post spread is strongly biased for the seller who is favored by the prior belief. Therefore, in equilibrium this seller becomes a "global monopolist" and covers the entire market by posting a limit price which excludes the competitor. The competitive pressure by the non-selling firm keeps the price low and leaves rents to both types of buyer. In summary:

Proposition 2 (Pooling equilibrium). *For $q > q^P(\alpha)$ the unique cautious equilibrium is pooling on good 1: both buyer types purchase from seller 1 and enjoy rents, prices are $P_0(q,\alpha) = 0, P_1(q,\alpha) = 2f_0(q,\alpha) - 1$, profits for the sellers $V_0(q,\alpha) = 0, V_1(q,\alpha) = 2f_0(q,\alpha) - 1$, and expected payoff of the buyer $V_B(q,\alpha) = 1 + q - 2f_0(q,\alpha)$.*

As made clear by this proposition and illustrated in Figure 3, the standard Bertrand outcome in an asymmetric setting is robust to the introduction of small amounts of private information, i.e. for $\alpha < (q^P)^{-1}(q)$. In the special case with an ex-ante perfectly informed buyer ($q = 1$), $P_0(1,\alpha) = 0, P_1(1,\alpha) = 1$ and $V_0(1,\alpha) = 0, V_1(1,\alpha) = 1$ and the buyer surplus is $V_B(1,\alpha) = 0$. In this case the superior seller monopolizes the market and extracts the full surplus of the buyer.

4.2.3. Mixed Strategy Equilibrium

In the classic Hotelling (1929) pricing game with uniform distribution of consumers, a pure-strategy equilibrium fails to exist when the sellers are located relatively close to each other (see e.g. d'Aspremont, Gabszewicz and Thisse (1979)). Osborne and Pitchik (1987)

[7] The non-selling firm in a pooling equilibrium could reduce the profit of the selling firm to any non-negative level by posting a negative price (a weakly dominated strategy). For $q > q^P(\alpha)$ there is a continuum of pure strategy pooling equilibria, where the non-selling firm posts $P_0 \in [1 - 2f_0, 0]$ and firm 1 sells at price $P_1 = 2f_0 - 1 + P_0 \in [0, 2f_0 - 1]$ thereby achieving a profit of $V_1 = P_1$. In all these equilibria the tie arises with probability one.

is the best attempt of characterization of the mixed-strategy equilibrium in the original Hotelling model. Similarly, in our model the equilibrium is in mixed-strategy for weak ex-post differentiation, corresponding to intermediate levels of horizontal differentiation, i.e. between the separating and the pooling regions $\max\langle 1/2, q^S(\alpha)\rangle < q < q^P(\alpha)$. In j we clarify the logic and the intuition behind the characterization of the equilibrium in our more tractable setting with discrete distribution of demand.

Searching directly for a mixed strategy equilibrium with continuous action spaces and proving its uniqueness appear daring enterprises, as uncountable partitions of the action space are potential supports of equilibrium strategies. In a sequence of preliminary steps (Lemmata 1-6 in the Appendix) we obtain joint restrictions on the forms of the equilibrium randomizations that make construction and verification of uniqueness much simpler. The main restrictions are now summarized. First, both sellers must make positive expected profits. Second, firm j may quote a price P_j with positive probability (an atom) if and only if the opponent firm $1-j$ has a gap in its support, which includes the price $P_{1-j} = IC_i(P_j)$ corresponding to the opponent's atom through one of the two IC constraints. Third, only the highest prices that the buyer may accept, $P_0 = 1 - f_0$ and $P_1 = f_1$, can be played with positive probability, so that ties happen with probability zero in equilibrium. Finally, almost all prices in the support of one firm's strategy must correspond through IC constraints to a single price played by the opponent. These results could be of independent interest and may shed some light on equilibrium play in other games with continuous action space and countable types of private information.[8]

Building on these properties, we first present a unified formulation of the equilibrium and then briefly describe the equilibrium strategies for each parameter configuration. As illustrated in Figure 3, the mixed strategy region in the parameter space can be partitioned in exactly four sub-regions (M1-M4), each corresponding to a different specification of the equilibrium strategies. We omit the minute details of the derivations, which are available on request.

General Characterization. Three benchmark prices help picture the form of seller j's equilibrium mixed strategy, whose c.d.f. is denoted by G_j, and the shape of its support, S_j. First, consider the *upper and lower bounds* to S_j. Seller $j = 0, 1$ randomizes on a set $S_j \subseteq [\underline{P}_j, \overline{P}_j]$, where, depending on parameters (q, α), these bounds take values in the following intervals: $\underline{P}_0 \in [0, \overline{P}_0), \overline{P}_0 \in (1-f_1, 1-f_0], \underline{P}_1 \in [2f_0 - 1, \overline{P}_1], \overline{P}_1 \in [2f_1 - 1, f_1]$.

Next, let $\tilde{P}_j \in S_j \subset [\underline{P}_j, \overline{P}_j]$ be the *fully separating* price, at which seller j sells with probability $\Pr(\sigma_j)$ to its own customer, given the strategy of the competitor. The following

[8] In the Appendix we explain the close connection with analogous restrictions derived by Osborne and Pitchik (1987) for the Hotelling model with uniform distribution (see their Appendix 1). Our additional results (Lemma 5 and 6) allow us to prove uniqueness of the equilibrium.

nonlinear map links this price to the upper and lower bound of the support.

$$p_j\left(\underline{P}_j,\overline{P}_j\right) := \frac{\Pr(\sigma_j)\cdot\underline{P}_j\cdot\overline{P}_j}{\Pr(\sigma_j)\cdot\overline{P}_j-\Pr(\sigma_{1-j})\cdot\underline{P}_j}.$$

Proposition 3 (Mixed strategy equilibrium). *For any $\alpha\in[0.5,1)$ and for any $q\in\left[q^S(\alpha),q^P(\alpha)\right]$ there is a mixed strategy equilibrium where, for each seller, upper bound and lower bound of the support and the fully separating price, $\{\overline{P}_j,\underline{P}_j,\tilde{P}_j\}_{j=0,1}$, are the unique solution to the non-linear system of equations*

$$\begin{array}{ll}\tilde{P}_0=\min\langle 1-f_0,p_0\left(\underline{P}_0,\overline{P}_0\right)\rangle & \tilde{P}_1=\min\langle f_1,p_1\left(\underline{P}_1,\overline{P}_1\right)\rangle\\ \overline{P}_0=\min\left\langle 1-f_0,1-2f_0+\tilde{P}_1\right\rangle & \overline{P}_1=\min\left\langle f_1,2f_1-1+\tilde{P}_0\right\rangle\\ \underline{P}_0=1-2f_1+\tilde{P}_1 & \underline{P}_1=2f_0-1+\tilde{P}_0\end{array}\tag{4.4}$$

which also satisfies

$$\underline{P}_0<\min\langle\alpha(1-q),1-f_1\rangle,\quad \underline{P}_0>\max\langle 0,1-2f_1+f_0\rangle.\tag{4.5}$$

The support of seller j's strategy S_j is the entire interval $\left[\underline{P}_j,\overline{P}_j\right]$ except (1) if the solution to (4.4) and (4.5) entails $\tilde{P}_0>1-f_1$, as in the M1 and M2 regimes (parameters regions), (resp. $\tilde{P}_1>f_0$, as in M1, M2, and M3) then S_0 (resp. S_1) does not include the interval $(1-f_1,\tilde{P}_0)$ (resp. $(f_0,\tilde{P}_1)$) and correspondingly seller 1's (resp. seller 0's) strategy has an atom on the maximum price f_1 (resp. $1-f_0$); (2) if $\tilde{P}_1=f_1$, as in M1, then $S_0=[\tilde{P}_0,1-f_0]$. The equilibrium payoff of seller j is $V_j=\Pr(\sigma_j)\tilde{P}_j$, the fully separating price times the probability of selling at that price.

Depending on the parameter couple (q,α) defining the game analyzed, the min functions in (4.4) select different combinations of their arguments, and the system (4.4) delivers a different type of solution, corresponding to a different form of equilibrium. Only four combination of (4.4) are consistent with no profitable deviations, and only one is possible for any given parameter pair. We now provide a verbal description and discussion of these four equilibrium regimes, each prevailing in the corresponding region of parameters as indicated in Figure 3. The strategies are illustrated in Figures 6 to 9. We have separately checked (4.4) and the restriction (4.5) for each candidate equilibrium. The Appendix illustrates the details of this tedious exercise only for M1, the other three being similarly constructed. The missing parts are available on request.

The Four Mixed Strategy Equilibrium Regimes: M1. A small increase of the prior belief q from the separating level $q^S(\alpha)$ — and similarly a reduction of the precision of the signal α — raises the posterior valuation for good 1 of the consumer with unfavorable signal σ_0, giving firm 1 an incentive to be more aggressive than in the separating region. In particular, the best reply of firm 1 jumps from $IC_1(P_0)$ to $IC_0(P_0)$ at the interior switching price $\hat{P}_0 = 1 - 2f_0 + \alpha q < 1 - f_0$. In the region of parameters M1 (Figure 3) firm 1's equilibrium strategy puts some weight on low prices in order to attract the type-0 customer with some probability. Firm 0 responds by posting corresponding low prices. Each seller still posts the highest feasible price, f_1 and $1-f_0$ respectively, with positive probability (an atom), and spreads the remaining probability with an atomless distribution on an interval of prices. As illustrated in Figure 6, consistently with the general restrictions mentioned earlier and derived in the Appendix, the probability mass by seller 1 on prices above f_0 consists *only* of an atom at $P_1 = f_1$ in region M1, and there is a gap in firm 1's support S_1 between $f_0 = IC_0(1 - f_0)$, recalling that on $P_0 = 1 - f_0$ firm 0 puts an atom, and f_1.

M2. By increasing q (or reducing α) beyond the boundary between regions M1 and M2 in Figure 3, the M1 equilibrium breaks down, since seller 0 would profit from deviating to price $1 - f_1$, thereby gaining the demand of the type-1 consumer when the opponent posts f_1. Intuitively, with even lower ex-post differentiation in customers' valuations, the incentive for seller 1 to separate types is reduced and the incentive to insist on low (pooling) prices is enhanced. The equilibrium of type M2 (Figure 7) is like M1, with the addition of the interval $[\widetilde{P}_1, f_1]$ to S_1, and of the corresponding prices $[\underline{P}_0 = IC_1(\widetilde{P}_1), 1 - f_1]$ to S_0. The probability mass by seller 1 on prices above f_0 consists not only of the atom on f_1 (as in M1) but also of the density on the interval $[\widetilde{P}_1, f_1]$. Finally, there is a hole in the support of each player corresponding to the atom by the competitor on the highest price.

M3. Increasing further the prior belief and/or decreasing further the signal precision, the effects illustrated in M2 are reinforced. Ex-post differentiation in the valuations becomes so low that the ex-ante superior seller 1 does not play at all the highest price f_1. In the mixed strategy equilibrium of type M3 (Figure 8) seller 0's strategy has no holes and an atom on $1 - f_0$; seller 1's strategy has a hole between $[\underline{P}_1, f_0]$ and $[\widetilde{P}_1, \overline{P}_1]$, with $\widetilde{P}_1 = 2f_1 - 1 + \underline{P}_0 > f_0$, and no atom.

M4. Finally, when the prior belief is rather balanced but the precision of the signal is very low (close to 1/2) players compete aggressively on customers who are mildly differentiated ex post. Ex-post differentiation is still strong enough to prevent a pooling equilibrium from arising on the ex ante superior good 1. In the M4 equilibrium (Figure 9) not even seller 0 plays the highest price $P_0 = 1 - f_0$, and there are no atoms in the equilibrium randomizations, nor holes in their supports, which are connected.

As we have seen, for $q > q^P(\alpha)$ ex post differentiation is so weak, vis-à-vis ex ante beliefs, that the game exhibits a pooling equilibrium: firm 1 clears the market by selling at a low, limit price, which prevents the competitor from selling above marginal cost, given the strong ex-ante disadvantage. Notice that M1 results either in a separating outcome or in a pooling outcome on good 1, while the other mixed equilibria result possibly in a separating or in a pooling outcome on either good. As parameters change from the separating to the pooling region, the equilibrium mixed strategies shift gradually probability mass away from high prices to more aggressive prices, leaving increasing rents to the buyer.

4.2.4. Uniqueness of Equilibrium

As proven in the Appendix, our characterization of equilibrium regimes is completed by:

Proposition 4 (Uniqueness of equilibrium). *The cautious equilibrium is unique for almost all parameters $\alpha \in [1/2, 1]$ and $q \in [1/2, 1]$.*

In the mixed strategies equilibrium regions, the explicit construction of the strategies (omitted) shows that the four regimes are mutually exclusive. Therefore, to establish uniqueness of the equilibrium we show that the mixed strategy equilibrium cannot take any form other than the four described above. The general restrictions on the equilibrium mixed strategies of the two sellers, also illustrated in the Appendix, make this task relatively simple. If no atoms are played, the support of each strategy is connected, and we are in the M4 regime. Conversely, to any atom played, necessarily on the highest IR prices f_1 and $1 - f_0$, there must correspond a gap in the support of the opponent. With atoms, we need to consider only a few alternative candidates for equilibrium, discarded directly in the proof.

4.2.5. Equilibrium with Ex-Ante Homogeneous Products

As an illustration and for future reference, we report the equilibrium for the special case of ex-ante identical sellers ($q = 1/2$). There is an interval of α corresponding to each of the three types of equilibria which are not asymmetric in nature (M4, M2 and separating). For $\alpha \in \left[1/2, 2 - \sqrt{2}\right]$ the unique equilibrium is of type M4, defined by $\underline{P}_0 = \underline{P}_1 = \sqrt{2 - 8\alpha + 8\alpha^2})$ and $\widetilde{P}_0 = \widetilde{P}_1 = \left(1 + \sqrt{2}\right)(2\alpha - 1)$ and with profit $V_0\left(1/2, \alpha\right) = V_1\left(1/2, \alpha\right) = \left(1 + \sqrt{2}\right)(2\alpha - 1)/2$ for both sellers. For $\alpha \in \left[2 - \sqrt{2}, 2/3\right]$ the unique equilibrium is of type M2, defined by $\underline{P} = 1 - 3\alpha - 2\alpha^2/(1 - 4\alpha + \sqrt{1 - 8\alpha + 12\alpha^2})$ and $\widetilde{P} = (2\alpha - 1 + \sqrt{1 - 8\alpha + 12\alpha^2})/2$ and with profit $V\left(1/2, \alpha\right) = (2\alpha - 1 + \sqrt{1 - 8\alpha + 12\alpha^2})/4$. Finally for $\alpha \in [2/3, 1]$ the separating equilibrium is $P = \alpha$, and $V\left(1/2, \alpha\right) = \alpha/2$ for both sellers.

5. Equilibrium Comparative Statics: Value of Public and Private Information

In this section we derive and discuss our main results. The complete characterization of the unique equilibrium allows us to analyze the dependence of the equilibrium payoffs of the buyer and the sellers on the degree of ex-ante heterogeneity of the sellers and the precision of the private information of the buyer. We focus mostly on two properties related to the equilibrium value of information for each party: convexity in the prior belief q, which pertains to the value of *public* information, and monotonicity in the precision of private information α, which indicates the value of contemporaneous *private* information of the buyer. Our results are compared with those obtained in the growing literature on the value of information and learning with endogenous pricing.

In this paper we compare situations where the quality of information acquired by the buyer is known or observed. This corresponds to the analysis of the second stage of price competition which follows a first stage where sellers and buyer interact to determine the information structure of the buyer. Our model could be extended by modelling the first stage of such a dynamic game. Typically, sellers affect the precision of the information of buyers by allowing them to try the product or to return it if not satisfied. In a similar fashion, the buyer could control the information acquired at a cost. Finally, an alternative to our comparative statics approach would be to analyze a game of covert information acquisition by the buyer where the amount of information is determined in equilibrium, as done by Crémer, Khalil and Rochet (1997) in a single-principal setting. The extension of our model with competing principals to the case of contemporaneous information acquisition is left for future research.

In order to reach a better understanding of the role of market power and competition, we first describe the efficient allocation in Section 5.1 and the solution of the problem of the monopolist in Section 5.2. We then return to our duopoly model, where we derive the value of information for the duopolists in Section 5.3 and for the buyer in Section 5.4.

5.1. Efficient Allocation

Given that production costs are set to zero, the total surplus to be divided among the three players is equal to the valuation of the buyer. In the efficient allocation the consumer buys from the ex-post superior seller: the consumer of type i chooses good 1 if $f_i \geq 1-f_i$, or $f_i \geq 1/2$, and good 0 otherwise. The social optimum can be easily implemented in this model by giving the bargaining power to the informed buyer. For $q \in [1/2, \alpha]$ the signal is relevant for the optimal decision, and the social surplus is $S(q,\alpha) = \Pr(\sigma_1) f_1 + \Pr(\sigma_0)(1 - f_0) = \alpha$. For $q \in [\alpha, 1]$ it is efficient to buy good 1 regardless of the signal, so that the buyer's ex ante valuation or expected social surplus is $S(q,\alpha) = \Pr(\sigma_1) f_1 + \Pr(\sigma_0) f_0 = q$. Overall social surplus is a piece-wise, continuous, weakly increasing and convex function of q for given α, and similarly a concave function of α for given q. Therefore, revelation of additional information introduces a spread in the belief which can only increase social welfare by allowing for better decisions. A more precise signal (higher α) leads to an increase in the total surplus only if it is strong enough to potentially reverse the prior; otherwise ex-post information is socially irrelevant. Similarly, a higher prior for one good does not help when the signal is strong enough to reverse it.

5.2. Monopoly Benchmark

Consider the simple optimization problem of a monopolist competing against a good sold at fixed price (set to zero for convenience). The monopolist can decide not to sell at all by posting a non-selling price $P^N > 2f_1 - 1$, sell only to type-1 buyer by posting the separating price $P^S = 2f_1 - 1$ with expected payoff $q - (1 - \alpha)$, and sell for certain by posting the pooling price $P^P = 2f_0 - 1$ with payoff $2f_0 - 1$. The optimal pricing strategy is

$$P = \begin{cases} P^N & \text{for } \; q \leq 1 - \alpha \\ P^S & \text{for } \; 1 - \alpha \leq q \leq q^P(\alpha) \\ P^P & \text{for } \; q \geq q^P(\alpha), \end{cases}$$

where $q^P(\alpha)$ is again the largest root of the quadratic equation (4.3). For low enough prior belief ($q < 1 - \alpha$) the monopolist prefers not to sell since in this region even the separating price is negative $P^S < 0$. For intermediate beliefs the separating price gives a higher expected payoff than the pooling price. For a high enough prior ($q \geq q^P$) pooling becomes optimal for a similar reason to that discussed in the Section 4.2.2.

The resulting monopolist profit function V_M is (strictly) increasing (when positive) and globally convex in q, being the maximum of convex functions:

$$V_M(q, \alpha) = \begin{cases} 0 & \text{for } \; q \leq 1 - \alpha \\ q - (1 - \alpha) & \text{for } \; 1 - \alpha \leq q \leq q^P \\ 2f_0(q, \alpha) - 1 & \text{for } \; q \geq q^P. \end{cases} \tag{5.1}$$

Convexity of the profit function of the monopolist in the prior distribution is a very general property with important implications regarding revelation of public information.[9] For instance, it *always* holds for a general number of signals in the binary state monopoly pricing model since the conditional signal distributions satisfy the monotone likelihood ratio property without loss of generality. More generally, it can be shown that the property holds for an affiliated environment. It is a manifestation of the linkage principle of Milgrom and Weber (1982), which implies that in a second-price auction with affiliated valuations

[9] See Ottaviani (1997) for more details on these results.

of the buyers revealing public information is always beneficial to the seller.[10] Monopoly pricing with a single seller can be seen as a second-price auction, with the monopoly price playing the role of the reserve price. The bid of the buyer is either above the reserve price, in which case the reserve price is the second-price paid by the buyer, or below it when the buyer decides not to buy.

Convexity of profits in the prior distribution implies that any change in the environment that spreads the posterior belief benefits the decision maker. In particular, a spread is achieved by revealing public information about the quality of the good, because of the martingale property of beliefs. The convexity property has also important dynamic consequences, for it implies that the monopolist values learning. A patient monopolist would then be willing to spend resources in the short run to foster revelation of public information on the quality of the good.

Finally, notice that monopoly profits are non-monotonic in the precision α of the buyer's signal. As it is seen immediately from (5.1), the monopolist's profit function is decreasing in α for $\alpha \leq (q^P)^{-1}(q)$, and increasing for $\alpha \geq (q^P)^{-1}(q)$.[11] The monopolist benefits from a stronger signal of the buyer when posting the separating price. In the pooling region the monopolist is instead forced to reduce further the price to sell to the buyer with a more precise unfavorable signal.

5.3. Value of Information for the Duopolists

We are now ready to discuss whether the properties of the profit function of the monopolist extend to a strategic setting.[12] First, without making use of the characterization of the mixed-strategy equilibrium we show below that convexity in the prior belief breaks down. Second, our characterization of the equilibrium in the mixed-strategy region delivers a surprising result of non-monotonicity of the duopolist's profit in the prior belief. We also briefly report on the properties of the sum of the equilibrium profits of the sellers, particularly important for dynamic extensions of the model. Finally, a seller's payoff is non-monotonic in the precision of the buyer's private information.

Value of Public Information We now establish our main result:

[10] This property has been used by Milgrom and Weber (1982) to rank the revenue produced by different auction formats, which entail different revelation of public information.

[11] The monopolist's solution is asymmetric with respect to $q = 1/2$. When $q < 1/2$ the monopolist's profits are increasing in α for $q \geq 1 - \alpha$ and constantly equal to 0 otherwise.

[12] At a more technical level, we remark that the equilibrium payoff of each duopolist is continuous in both parameters. Continuity is easily verified algebraically at the borders across the regions. The *form* of the equilibrium randomizations needs not instead be continuous when crossing such borders between regions. For instance, as we cross the boundary between the $M1$ and $M2$ regions, the atom at the upper bound of S_1 changes discontinuously. Nonetheless it is verified that in such cases there are two equilibria, so that the equilibrium correspondence is upper-hemicontinuous in the parameters.

Proposition 5 (Non-convexity of seller's profit in prior belief). *For high enough quality of private information $\alpha \geq 2/3$, the equilibrium profit function of firm 1 is not convex in the prior belief q.*

This is easily proved by projecting the linear segment αq (which corresponds to the payoff in the separating equilibrium region, nonempty for $\alpha \geq 2/3$) to the border of the pooling region $q = q^P$. The payoff cannot be convex because αq^P is strictly larger than the equilibrium payoff $V_1(q^P, \alpha) = q^P - (1-\alpha)$ computed at the same belief. See Figure 10 for an illustration.

The non-convexity originates from the high level of profits enjoyed by the duopolists in the separating equilibrium, resulting when the prior belief is not too biased in favor of one particular product. Intuitively, as the prior inclination of the market for good 1 rises from the separating to the pooling region, and the ex-post valuations of both buyer's types with it, the total payoff of seller 1 rises. But this payoff increase occurs at a less than proportional rate with q, because seller 1 goes from fully extracting the buyer's rent in the separating region to a limit price which leaves some rents to both types of buyer in the pooling region. As discussed in Section 3, for $\alpha > 1/2$ an increase in q augments vertical differentiation by increasing both f_1 and f_0, but also decreases horizontal differentiation by reducing the spread in the distribution of valuations Δf for $q > 1/2$. Lower horizontal differentiation increases the aggressiveness of the sellers and results in lower prices. In particular, $\tilde{P}_1$ (as well as the other bounds of S_1) decreases in q. The value function $V_1 = \Pr(\sigma_1)\tilde{P}_1$ increases less than linearly in q in M2 and M3, because $\Pr(\sigma_1)$ increases linearly in q, while $\tilde{P}_1$ decreases.

This non-convexity implies that seller 1 — favored by the prior belief — strictly prefers to avoid diffusion of public information to customers for intermediate values of the prior belief, even if this information is more likely to make the customer more willing to buy good 1. The duopolists are hurt by more information on relative quality, because it results in a more asymmetric situation and therefore more pooling in equilibrium. The comparison with the monopoly case reveals that the strictly negative value of learning can only be due to purely strategic considerations. Notice also the contrast to cases of public learning, studied for instance by Bergemann and Valimaki (1997), where instead the equilibrium value functions of the sellers are convex. The linkage principle breaks down when a second price-setting seller who offers a differentiated good is added to the picture. An important implication of this for auction theory is that the revenue ranking of Milgrom and Weber (1982) does not extend to settings where the sellers are competing in mechanisms.

The same argument used for the single seller shows that the *sum* of the equilibrium profits of both sellers is also not convex in the prior belief when the quality of the buyer's information is high enough. Finally, is the sum of all three players expected equilibrium payoffs convex in the prior, like the payoff achieved in the social optimum? In the pooling region, the equilibrium is efficient and the sum of players' payoffs is q. In the separating region, the firms extract full rents from the buyer, and the sum of payoffs is α. The total surplus in equilibrium is continuous in q, but the separating equilibrium extends beyond the efficient boundary $q^S(\alpha) = \alpha$. Since for intermediate values of the prior belief q, a public and informative signal would lead with some chance to an inefficient equilibrium, where some surplus is wasted, it can lead to lower social surplus. For high enough quality of private information the social surplus achieved in equilibrium is not convex in q. The result is proved again by projecting the linear segment q, equal to the total surplus in the pooling region, to $q = q^S(\alpha) > \alpha$ (for $\alpha > \sqrt{2}/2$), where α is the total (separating) payoff.

Non-Monotonicity in Prior Belief. The profit function of the monopolist (selling good 1) is (strictly) increasing (when positive) in q. A monopolist with better prior quality achieves higher profits. Surprisingly, this is not necessarily true in the duopoly model. As q rises, the temptation for seller 1 of clearing the market becomes stronger, and triggers an aggressive response by the competitor. It can be shown by simple algebra that for given α the equilibrium profit of seller 1, V_1, is always strictly larger at the M1-M2 boundary q^{M1} than at the M3-Pooling boundary q^P, and therefore is necessarily decreasing in q in M2 and/or M3. For concreteness, Figure 10 displays V_1 and V_0 as functions of q for $\alpha = .69$; in region M3 $V_1 = \Pr(\sigma_1)\tilde{P}_1$ is decreasing in q, as the decline in $\tilde{P}_1$ due to fiercer competition is so strong to dominate the linear increase in $\Pr(\sigma_1)$. The ex-ante superior seller achieves lower profits despite being more favored by the prior belief.

Next consider the *sum* of the payoffs of the two sellers. In the model of Shaked and Sutton (1982) an increase in the level of vertical differentiation leads to less competition and more profits for the sellers. Does this prediction extend to our model where differentiation is the result of private information on the relative quality of the different competitors? Without private information ($\alpha = 1/2$) we have $V_0(q, 1/2) + V_1(q, 1/2) = \max\langle 1 - q, q\rangle$. For $q = 1/2$ both sellers are equally good for the buyer and make zero profits, while an asymmetric belief results in higher profits of the superior seller and constant zero profits for the inferior seller. While in the case without private information the more vertically differentiated the goods, the higher the sum of the sellers' profits, this is no longer true in the presence of private information. On the one hand, the increased level of vertical differentiation leads to more profits for the sellers, along the lines of Shaked and Sutton (1982). On the other hand, the induced lower horizontal differentiation increases the competitive pressure on prices and tends to lead to lower profits for the sellers. The latter effect dominates the former for intermediate prior beliefs (relative to private information), while the opposite is true for extreme priors. The sum of the equilibrium profits of seller i is non-monotonic in the level of vertical differentiation $|q - 1/2|$ for any level of noisy private information $\alpha \in (1/2, 1)$. The result is proved easily by comparing the sum of the equilibrium profits at $q = 1/2$ and $q = q^P(\alpha)$. Notice that $\sum_{j=0}^{1} V_j(q^P(\alpha), \alpha) = q^P(\alpha) - (1 - \alpha)$, and $\sum_{j=0}^{1} V_j(1/2, \alpha) = 2V(1/2, \alpha)$ reported in Section 4.2.5. By direct comparison the first quantity is strictly larger than the second for any α, so that the sum must be decreasing in q in part of the interval $[1/2, q^P(\alpha)]$. The sum of the profits is instead strictly increasing in q for $q \geq q^P(\alpha)$, being equal to the pooling profits of the ex ante superior seller 1.

Value of Contemporaneous Private Information It can be immediately verified from the closed-form solutions given in Section 4.2.5 that for $q = 1/2$ the value function of a duopolist is monotonic and concave in the quality of private information α of the buyer. As soon as the prior belief becomes asymmetric, this monotonicity does not hold any more, similarly to what happens in the solution to the monopoly model. $V_1(q, .)$ is decreasing in α in the pooling region, i.e. for $\alpha \leq (q^P)^{-1}(q)$, and increasing in the separating region, i.e. for $\alpha \geq (q^S)^{-1}(q)$. In the separating region the payoff of the sellers increases in the amount of private information of the buyer, in accordance with the principle that, as quality becomes more different, price competition between less similar products leads to an increase in the profits of either seller.[13] This non-monotonicity originates from endogenous pricing, but it is not due to strategic reasons because it is also present in the monopoly case.

Proposition 6 (Non-monotonicity of seller's profit in private information). *The equilibrium profits of seller 1 are decreasing in the precision α of the signal of the buyer for $\alpha \leq (q^P)^{-1}(q)$ and increasing for $\alpha \geq (q^S)^{-1}(q)$.*

5.4. Value of Information for the Buyer

It is natural to consider what would happen if the consumer decided how much private information to acquire on the quality of the products. This corresponds to the Bertrand game between the sellers being played after observing the private signal's quality α chosen by the buyer. The main result is that the buyer benefits from small amounts of private information, but is hurt by too precise a private signal, which reduces the competitive pressure on prices. This contrasts with the well-known fact that more information necessarily benefits in ex ante terms a decision maker facing fixed prices, as well as in our game the society of buyer and sellers. A negative value of information is not a new result in a game-theoretic context, but this model clarifies very sharply the nature of the strategic interaction behind it.

[13] On this see e.g. Shaked and Sutton (1982).

In the pooling equilibrium region of parameters, the buyer's expected payoff $V_B^P(q,\alpha) = 1+q-2f_0(q,\alpha)$ is increasing in the quality of private information α. In the pooling region, the seller is forced to reduce further the price to attract the buyer with a stronger negative signal. For the same reason, pooling becomes increasingly costly as α rises until the temptation to clear the market loses its appeal. This effect is eventually reversed because more precise information corresponds to higher ex-post differentiation of valuations, and therefore to more rent-extraction by the sellers in part of the mixed strategy region and in the separating regime. The optimal amount of costless information acquired by the buyer is interior to the mixed strategy region. Once prices are endogenous, the buyer would sometime prefer not to gather information even if it were free:

Proposition 7 (Non-monotonicity of buyer's surplus in private information). *The buyer benefits from an informative private signal of low precision α, but too precise signals lead to lower payoff in equilibrium.*

As of convexity in prior beliefs q, it is easy to show that it does not hold for the buyer either. In the pooling equilibrium, the buyer obtains $q-(2f_0-1)$ from good 1, strictly concave in q. When buying the same good, the buyer obtains the expected posterior valuation for good 1, equal to the prior belief q, and pays the pooling price which is convex in the prior since ex post differentiation fades away with the rising prior.

6. Robustness

6.1. Sellers Competing in Mechanisms

Consider the possibility of expanding the tools given to the competitors by allowing the sellers to offer menus of lotteries among which the buyer can choose. In this section we show that excluding menus of contracts — as done in the analysis so far — is without loss of generality in our model. This is a non-surprising implication of the fact that the payoff of our (risk-neutral) buyer is linear in the valuations for the (indivisible) good. Nevertheless, our analysis is restrictive in the sense that we do not allow the sellers to condition the offer in their mechanism on the offer of the competitor. See Epstein and Peters (1998) for an investigation of the issues involved when such a dependence is allowed.

A *lottery contract* $\{\lambda_{kj}, T_{kj}\}$ offered by seller j consists of a probability λ_{kj} of transferring the good to the buyer who accepts the lottery, and an unconditional transfer T_{kj} from that buyer to the seller. A *menu* of lottery contracts is a collection of lottery contracts. A menu is *degenerate* if has at most one distinct lottery which allocates the good to the buyer with positive probability. A lottery k with $\lambda_{k.}=1$ is degenerate, being equivalent to a simple price offer.

To the purpose of characterizing the optimal menu of seller 1 in response to a menu of lotteries offered by the opponent 0, it is convenient to define the outside option of the consumer with signal σ_i as $u_i = \max \langle \max_k \lambda_{k0}(1-f_i) - T_{k0}, 0\rangle$, the highest payoff achieved by either accepting the most preferred lottery offered by seller 0 or rejecting all lotteries. Notice that $1 - f_0 > 1 - f_1$ implies $u_0 \geq u_1$.[14] The problem for seller 1 is then a standard mechanism design problem with type-dependent outside options. The menu of lotteries offered by seller 1 solves

$$\max_{\substack{\lambda_0,\lambda_1 \in [0,1] \\ T_0, T_1}} \Pr(\sigma_0)T_0 + \Pr(\sigma_1)T_1$$

s.t.

$$\lambda_0(\lambda_0 f_0 - T_0 - u_0) \geq 0 \qquad (PC_0)$$

$$\lambda_0\lambda_1[\lambda_0 f_0 - T_0 - (\lambda_1 f_0 - T_1)] \geq 0 \qquad (SS_0)$$

$$\lambda_1(\lambda_1 f_1 - T_1 - u_1) \geq 0 \qquad (PC_1)$$

$$\lambda_0\lambda_1[\lambda_1 f_1 - T_1 - (\lambda_0 f_1 - T_0)] \geq 0 \qquad (SS_1)$$

where the second index of the lotteries referring to seller 1 has been dropped. According to the participation constraint PC_i the buyer of type i weakly prefers to buy from seller 1 — requiring both $IR_{i,1}$ and $IC_{i,1}$ to be satisfied — whenever the seller does not exclude that type of buyer. The self-selection constraint SS_i requires type i buyer to prefer lottery $\{\lambda_i, T_i\}$ to $\{\lambda_{1-i}, T_{1-i}\}$.

We will now show that the best reply of a seller to any menu of lotteries of the opponent is a degenerate menu consisting of a simple price offer to the buyer, so that there cannot be an equilibrium in non-degenerate menus. In particular, when the type 0 buyer is excluded, the optimal menu consists only of the (separating) lottery $\{1, f_1 - u_1\}$, while when no type is excluded, the only distinct lottery in seller's 1 optimal menu consists of the (pooling) lottery $\{1, f_0 - u_0\}$. The prior q which makes the seller indifferent between separating and pooling satisfies

$$\Pr(\sigma_1|q)[f_1(q,\alpha) - u_1] = f_0(q,\alpha) - u_0. \qquad (6.1)$$

Notice that when also the opponent makes a simple price offer, so that $u_i = \max\langle 1 - f_i + P_0, 0\rangle$, (6.1) determines the switching price. The proof (in the Appendix) adapts standard arguments from non-linear pricing.

Proposition 8 (Degenerate menus). *There are no equilibria in non-degenerate menus.*

[14] The same argument clearly applies to mixed strategy over menus of lotteries.

6.2. A Model with a Continuum of Signals

The binary signal structure studied so far is admittedly restrictive and gives rise to non-trivial analytical problems. Nevertheless, it has allowed us to provide a complete illustration of the dependence of the equilibrium on the degree of ex-ante heterogeneity and on the amount of private information. In this section we show that alternative continuous formulations would not simplify these tasks. In the process, we also discover that some of our results are robust to more general specifications of signal distributions. To this purpose, we present a binary-state version of the model with continuous signals distributed with monotone (without loss of generality) and bounded likelihood ratio. In general, a pooling equilibrium results for extreme prior beliefs. We cannot say anything more without specializing the analysis to particular distributions. In all the examples that we have explored of continuous signals distributed symmetrically, equilibrium characterization and comparative statics are less tractable than in our discrete setting. By contrast, in our discrete setting we were able to describe a unique equilibrium for all parameter configurations.

Setup. We modify only the signal structure, leaving the rest of the model unchanged. Signals have distribution H_i (and density h_i) conditional on state ω_i on the same support $S \subseteq \mathbb{R}$, to rule out the case of shifting support. They are imperfectly informative, with bounded (though increasing) likelihood ratio $\lambda(\sigma) = h_1(\sigma)/h_0(\sigma) \in [\underline{\lambda}, \overline{\lambda}]$ for $0 < \underline{\lambda} < \overline{\lambda} < \infty$. The posterior belief of state ω_1 after signal σ is realized is:

$$f_\sigma(q) = \frac{h_1(\sigma)\,q}{h_1(\sigma)\,q + h_0(\sigma)\,(1-q)} = \frac{\lambda(\sigma)\,q/(1-q)}{1+\lambda(\sigma)\,q/(1-q)}.$$

The decision of the consumer with signal σ is determined by: the $IR_{\sigma,1}$ constraint, $f_\sigma(q) - P_1 \geq 0$; $IC_{\sigma,1}$, which is $f_\sigma(q) - P_1 \geq 1 - f_\sigma(q) - P_0$; $IR_{\sigma,0}$, $1 - f_\sigma(q) - P_0 \geq 0$; and $IC_{\sigma,0}$, $f_\sigma(q) - P_1 \leq 1 - f_\sigma(q) - P_0$. The highest IR prices for any buyer type are: $\overline{P}_1 = (\overline{\lambda}q)/(\overline{\lambda}q + 1 - q)$, and $\overline{P}_0 = (1-q)/(\underline{\lambda}q + 1 - q)$, both in $(0,1)$. Since $\lambda(\cdot)$ is increasing, $f_\sigma(q) \gtreqless x \Leftrightarrow \sigma \gtreqless \lambda^{-1}[x(1-q)/q(1-x)]$. Substituting $x = (1 + P_1 - P_0)/2 = (1+\Delta P)/2$ we obtain the cutoff signal satisfying both $IC_{\sigma,0}$ and $IC_{\sigma,1}$ with equality: $\sigma^{IC}(P_1 - P_0) = \lambda^{-1}[(1+\Delta P)(1-q)/(1-\Delta P)\,q]$. For all signal realizations $\sigma > \sigma^{IC}$ good 1 is strictly preferred to good 0, while the opposite is true for $\sigma < \sigma^{IC}$. Similarly, $x = P_1$ gives the cutoff signal $\sigma^{IR_1}(P_1) = \lambda^{-1}[P_1(1-q)/q(1-P_1)]$ corresponding to $IR_{\sigma,1}$. Clearly, $IR_{\sigma,1}$ is slack for all $\sigma > \sigma^{IR_1}$. Finally, $x = 1 - P_0$ gives $\sigma^{IR_0}(P_0) = \lambda^{-1}[(1-P_0)(1-q)/P_0 q]$.

To see which constraint binds, notice that for $P_1 + P_0 \leq 1$

$$\frac{P_1}{1-P_1} \leq \frac{1+\Delta P}{1-\Delta P} \leq \frac{1-P_0}{P_0}, \tag{6.2}$$

corresponding to $\sigma^{IR_1}(P_1) \le \sigma^{IC}(P_1 - P_0) \le \sigma^{IR_0}(P_0)$ since $\lambda^{-1}(.)$ is increasing, so that the IC constraint is binding for all σ (we drop σ here). Similarly to what happens in the binary case, if the price posted by firm 0 is low enough that $P_0 \le 1 - P_1$, then IC is binding and IR_1 slack, because $\sigma^{IC} \ge \sigma^{IR_1}$ and because both the IR_1 (i.e. $\sigma > \sigma^{IR_1}$) and the IC (i.e. $\sigma > \sigma^{IC}$) constraints need to be satisfied for a consumer with signal realization σ to buy good 1. When instead $P_1 + P_0 \ge 1$ the inequalities in (6.2) are reversed and $\sigma^{IR_0}(P_0) \le \sigma^{IC}(P_1 - P_0) \le \sigma^{IR_1}(P_1)$, so that the two $IR's$ are binding.

Best Reply. Let $E_q H(\sigma) = qH_1(\sigma) + (1-q)H_0(\sigma)$ denote the expected probability of receiving a signal no smaller than σ when the prior belief is q, and $\pi_1^{IC}(x, P_0) = x\left[1 - E_q H\left(\sigma^{IC}(x - P_0)\right)\right]$ and $\pi_1^{IR}(x) = x\left[1 - E_q H\left(\sigma^{IR_1}(y)\right)\right]$ the expected payoff to seller 1 from quoting price x when $IC_{\sigma,1}$ and $IR_{\sigma,1}$ binds, respectively. Next let:

$$P_1^{IC}(P_0) = \underset{0 \le x \le \min\left(1-P_0, \bar{P}_1\right)}{\arg\max} \pi_1^{IC}(x, P_0); \quad P_1^{IR} = \underset{1-P_0 \le y \le \bar{P}_1}{\arg\max} \pi_1^{IR}(y). \tag{6.3}$$

Then the best response of firm 1 is $P_1(P_0) = P_1^{IC}(P_0)$ if $\pi^{IC}\left(P_1^{IC}(P_0), P_0\right) > \pi^{IR}\left(P_1^{IR}\right)$ and $P_1(P_0) = P_1^{IR}$, independent of P_0, if $\pi^{IC}\left(P_1^{IC}(P_0), P_0\right) \le \pi^{IR}\left(P_1^{IR}\right)$. A similar construction holds for firm 0.

By the Theorem of the Maximum, the best reply map $P_j(P_{1-j})$ is a non-empty, upper hemi-continuous and compact-valued correspondence. It is easily proved that $P_1(P_0) > 0$ if and only if both $q > 0$ and $P_0 \in (\underline{P}_0, \bar{P}_0]$, where $\underline{P}_0 \equiv \left(1 - q - \overline{\lambda}q\right) / \left(1 - q + \overline{\lambda}q\right) < \bar{P}_0$. Symmetrically, $P_0(P_1) > 0$ if and only if both $q < 1$ and $P_1 \in [\underline{P}_1, \overline{P}_1]$, with $\underline{P}_1 \equiv (q - 1 + \underline{\lambda}q)/(1 - q + \underline{\lambda}q) < \overline{P}_1$. As long as the opponent quotes prices not above $\underline{P}_{1-j}$, firm j has no incentive to quote a positive price and opts out of the market.

Equilibrium. A Bayes-Nash equilibrium is a pair (P_0, P_1) such that $P_j = P_j(P_{1-j}(P_j)) \in [0, \overline{P}_j]$ for $j = 0, 1$. A *separating equilibrium* has $P_j > 0$ for $j = 0, 1$. It is *market-clearing* if $P_j + P_{1-j} \le 1$ and *rationing* otherwise. A *pooling equilibrium* on good j has $P_j > 0 = P_{1-j}$. In general:

Proposition 9 (Pooling for extreme prior beliefs). *For prior beliefs sufficiently favorable to seller* 1, $q \in [\overline{q}, 1]$ *with* $\overline{q} < 1$, *there exists a pooling equilibrium on good 1.*

Proof. If $P_0 = 0$ the IC constraint always binds: $\sigma^{IC}(\underline{P}_1 - 0) = \lambda^{-1}(\underline{\lambda}) = \underline{\sigma}$. Then:

$$\left.\frac{\partial \pi_1^{IC}(P_1 - P_0)}{\partial P_1}\right|_{P_0=0, P_1=\underline{P}_1} = 1 - \frac{(\underline{\lambda}q)^2 - (1-q)^2}{2q(1-q)\lambda'(\underline{\sigma})}\left[qh_1(\underline{\sigma}) + (1-q)h_0(\underline{\sigma})\right].$$

As $q \uparrow 1 > 1/(1 + \underline{\lambda})$, we observe that $(\underline{\lambda}q)^2 - (1-q)^2 > 0$ so that this quantity tends continuously to $-\infty$, and that $\overline{P}_1 \downarrow \underline{P}_1$. By continuity, for q close enough to 1, the slope of the objective function is strictly negative for *all* $P_1 \in [\underline{P}_1, \overline{P}_1]$. Therefore it is best to respond to $P_0 = 0$ with a price as small as possible, or $P_1(0) \leq \underline{P}_1$. By definition of $\underline{P}_1$, this implies that in turn $P_0 = 0$ is a best response to $P_1(0)$, and the claim follows: there is a pooling equilibrium on good 1, in which firm 0 posts a zero price and yet sells with probability zero.♦

Nothing more can be said in general. Specializing the analysis to even very simple signal densities does not help either. The comparison between maximized values in (6.3) quickly reaches forbidding peaks of algebraic complexity, which complicates the equilibrium characterization. We conclude that our discrete setting is preferable for the purpose of performing comparative statics across equilibria.

7. Conclusion

We have investigated price competition in markets where quality is difficult to ascertain and the price-taking buyer has private information on the relative quality of the alternative competitors. When the prior belief is very biased toward one good and private signals are not too informative, sellers compete fiercely and leave rents to the buyer as in the classical Bertrand model. When instead the prior is balanced and signals are of bounded but strong precision, the sellers become local monopolists. The more spread the distribution of valuations, due to a vaguer prior and stronger signals, the more rent the sellers can extract from the consumer. This fact has two main implications: First, contrary to the linkage principle of Milgrom and Weber (1982) the equilibrium profits of the seller are not globally convex in the prior belief, so that the sellers may lose from the release of public information. Second, the buyer is hurt by receiving too precise private signals, as long as the sellers know it.

Our insights on the value and incentives for information acquisition with strategic pricing can be applied naturally to a broad class of markets where buyers are offered individualized prices. The stylized relationship between a buyer and two competing sellers could be easily enriched in order to consider problems arising in the labor, credit, and insurance markets. It is essential that the price-taking party has superior information on the relative desirability of the competing price-setters. As a labor market application with the role of buyer and seller reversed, consider the situation of a job applicant (seller) with private information on the net costs of working for different employers (buyers) who compete in wage offers.

Finally, our static model is a building block for dynamic models of strategic pricing with private learning. For a dynamic extension of this model we refer to Moscarini and Ottaviani (1997)'s model of social learning about product quality with endogenous prices.[15]

[15] Related to this, Caminal and Vives (1996) construct a two-period model of duopoly competition for

The demand side of the market consists of a sequence of privately informed customers with the same preferences. Buyers are then able to partially infer the information possessed by other buyers by observing their purchase decisions as in the social learning model of Bikhchandani, Hirshleifer and Welch (1992). On the supply side of the market the sellers engage in repeated price competition. In this context prices serve not only the usual allocative role, but also act as a screening device for the transmission of the private information held by the previous buyers.

Appendix

A.I. General Properties of the Mixed Strategy Equilibrium

This Appendix illustrates general restrictions on the form of equilibrium mixed strategies for our game. All are valid for and only for the region of parameters where there is no pure strategy equilibrium, $q \in [\max\langle 1/2, q^S(\alpha)\rangle, q^P(\alpha)]$. The graphical representation of the constraints in the price space (Figure 2a) is of the great help to follow the arguments. A *mixed strategy* by firm j, $j = 0, 1$ is a probability measure over the Borel sets, and the *support* S_j of the mixed strategy is the smallest Borel set of probability 1.

Lemma 1 (Sellers' payoffs). *Both firms make strictly positive profits in equilibrium.*

Proof. Suppose that firm j makes zero profits in a mixed strategy equilibrium. If firm $1-j$ plays a positive mass of probability on prices above $IC_j(0)$, then firm j could post $P_j' = \varepsilon$ for some $\varepsilon > 0$, which would sell with positive probability, thereby making positive profits in contradiction with the assumption. Otherwise, we are back to the pooling equilibrium on good $1-j$, which does not exist in this region of parameters.♦

Recall that $\underline{P}_j$ and $\overline{P}_j$ denote the lower and upper bounds of the support of the equilibrium randomization of player j. Since, by Lemma 1, any price in the support - including the upper bound $\overline{P}_j$ - must yield positive expected profits, we have:

Lemma 2 (Mass above $IC_j(\overline{P}_j)$). *Seller $1-j$ must play prices above $IC_j(\overline{P}_j)$ with positive probability.*

Lemma 3 (Atoms and gaps). ***(i) Atoms and gaps must correspond through IC constraints.*** *If in equilibrium seller j plays a price P_j with positive probability (an atom), then there is a corresponding gap in the support of the opponent's randomization: $\exists G > 0$ such that $(IC_i(P_j), IC_i(P_j)+G) \cap S_{1-j} = \emptyset$ for $i = 0, 1$. The converse is also true, provided that seller $1-j$ plays prices weakly below the gap (some $P_{1-j} \leq IC_i(P_j)$).*
(ii) Atoms only at maximum prices. *Only the maximum prices that the buyer may accept, $P_0 = 1 - f_0$ and $P_1 = f_1$, can be played with positive probability by the sellers.*

a continuum of informed consumers.

Proof. (i) First, we show sufficiency of an atom for a corresponding gap. The price $IC_i(P_j) + \varepsilon$ for some $\varepsilon > 0$ is strictly dominated for firm $1-j$ by $IC_i(P_j) - \delta$ for some $\delta > 0$, because the latter steals a discrete mass of demand (the atom) from the competitor and loses only a small $\delta + \varepsilon$ in terms of unit revenues. Therefore $IC_i(P_j) + \varepsilon$ cannot possibly be in S_{1-j}, for a set of $\varepsilon \in (0, G)$, with $G > 0$ being the width of the gap.

Next, we show sufficiency of a gap $(IC_i(P_j), IC_i(P_j) + G)$ in S_{1-j}, provided further that prices weakly below $IC_i(P_j)$ are also in S_{1-j}, for a corresponding atom by firm j on P_j. First, prices in $(P_j, P_j + G)$ cannot be in S_j, since they are dominated by $P_j + G$, so the two gaps correspond through IC_i. By contradiction, suppose there is no atom on P_j. Then firm $1-j$ would gain strictly from playing $IC_i(P_j) + G$ rather than $IC_i(P_j) - \varepsilon$ for $\varepsilon \geq 0$ small, because $(P_j, P_j + G)$ are not played by firm j; but this contradicts the assumption that prices $IC_i(P_j) - \varepsilon$ for some $\varepsilon \geq 0$ small are in S_{1-j}.

(ii) Suppose that there is an atom at an interior price, e.g. at $P_1 < f_1$. Then, by (i) there is a corresponding gap in S_0 containing either $IC_0(P_1)$ or $IC_1(P_1)$ (or both), and firm 0 does not play prices in $(IC_i(P_1), IC_i(P_1) + \varepsilon)$ for some $\varepsilon > 0$, for either i. But then firm 1 would gain over P_1 by deviating to a strictly higher price $P_1 + \varepsilon$, which would sell with the same probability — positive by Lemma 1 — as P_1.♦

It follows immediately from claim (i) that, when there is an atom on $P_0 = 1 - f_0$ ($P_1 = f_1$), there must be a gap in S_1 containing $P_1 = f_0$ (resp. in S_0 containing $P_0 = 1 - f_1$). From claim (ii), it follows that each seller's equilibrium randomization G_i is continuous and, being nondecreasing by definition, has a density $g_i = G_i'$ a.e. for prices smaller than the maximum ones acceptable by the buyer. Finally, since $1 - f_0 \neq IC_i(f_1)$ for $i = 0, 1$:

Corollary 1. *Ties happen with probability zero in equilibrium.*

The previous results imply that the support of an equilibrium mixed strategy - a Borel set in $[0, 1]$, thus a countable union of bounded intervals - is a collection of *non-degenerate* intervals, plus possibly the upper bound of the support. In fact, absent any atom on interior prices, we may exclude any isolated point other than the maximum feasible price by considering the smallest set of prices played by a firm with probability one. The next result is that the "holes" separating these intervals must be projections through *one* IC constraint of the holes in the opponent's support.

Lemma 4 (Corresponding bounds). *The bounds of the disjoint intervals of prices that form the support of a player's equilibrium randomization must correspond through an IC_i constraints to those of the other player.*

Proof. By contradiction. Let $\check{P}_j$ be a lower bound of one of these intervals, such that $\check{P}_{1-j} \neq IC_i(\check{P}_j)$ for both $i = 0, 1$. Consider the case $\check{P}_{1-j} > IC_i(\check{P}_j)$. Then $P_j = IC_i(\check{P}_{1-j})$

dominates all prices in $(\check{P}_j, IC_i(\check{P}_{1-j}))$, in contradiction with the definition of equilibrium. Similarly if $\check{P}_{1-j} < IC_i(\check{P}_j)$.♦

These results are closely connected to analogous restrictions derived by Osborne and Pitchik (1987) in their analysis of the Hotelling pricing game with linear demand. They first try to solve the equilibrium of the pricing game for all possible pairs of firms' locations, which are their parameters. In their Appendix 1, they prove the claim of Lemma 1, their claim (b) is similar to our Corollary 1, claim (j) to Lemma 3, (ii), claim (m) to Lemma 3, (i). As a consequence, their partial characterization of different equilibrium regimes (T1 and T2) resembles ours (M1-M4). However, they cannot obtain enough restrictions to pin down uniquely the equilibrium for all parameter values. In our discrete setting we are able to do this.

Consider a price P_0 in the support S_0; by Lemma 4 there must be a price in S_1 corresponding to P_0 through one of the two IC constraints. The next result shows that there is almost always *only one* such a price: both the high price $IC_1(P_0)$ and the low price $IC_0(P_0)$ can be in S_1 only for a countable set of prices P_0; similarly for firm 0. Intuitively, the rate at which expected profits are lost by rasing a price, given the opponent's strategy, is different on the two IC constraints. This result simplifies enormously the search for mixed strategy equilibria over non countable action spaces.

Lemma 5 (The tie principle). *In equilibrium, for all values of the parameters (q, α), the set of prices in the support of seller j such that the two tying prices are both in the support of $1-j$, $\{P_j \in S_j : P'_{1-j} = IC_1(P_j) \in S_{1-j} \text{ and } P''_{1-j} = IC_0(P_j) \in S_{1-j}\}$, has Lebesgue measure zero.*

Proof. Consider seller 0 and tying prices by firm 1, the other case being symmetric. Fix any price P_0 in the relevant range $[0, 1-f_1]$, where two feasible (IR) tying prices by firm 1 exist: $P'_1 = 2f_1 - 1 + P_0$ and $P''_1 = 2f_0 - 1 + P_0$. Let us first compute the expected payoffs associated with these two prices, recalling that G_j denotes the c.d.f. of the mixed strategy played by seller j. P'_1 sells to both customers' types if $P_0 \geq P'_1 + 1 - 2f_0 = P_0 + 2\Delta f$, with chance $1 - G_0(P_0 + 2\Delta f)$; only to type 1 if $P_0 \in [P'_1 + 1 - 2f_1, P'_1 + 1 - 2f_0] = [P_0, P_0 + 2\Delta f]$, chance $G_0(P_0) - G_0(P_0 + 2\Delta f)$; to none otherwise. Thus the expected payoff to firm 1 from P'_1 is $\pi_1(P'_1) = \pi_{1,1}(P_0) = (2f_1 - 1 + P_0)[1 - \Pr(\sigma_0) G_0(P_0 + 2\Delta f) - \Pr(\sigma_1) G_0(P_0)]$. and, similarly, from P''_1: $\pi_1(P''_1) = \pi_{1,0}(P_0) = (2f_0 - 1 + P_0)[1 - \Pr(\sigma_0) G_0(P_0) - G_0(P_0 - 2\Delta f)]$

Contrary to the claim, suppose there exists a non-zero Lebesgue measure set X_0 of $P_0 \in S_0$ such that $P'_1, P''_1 \in S_1$. Since $P_0 \leq 1 - f_1 < 1 - f_0$, G_0 has a density $g_0 = G'_0$ at almost all points we are considering. $P'_1 \in S_1$ and $P''_1 \in S_1$ imply, for *all* $P_0 \in X_0$, that the two expected payoffs to firm 1 from P'_1, $\pi_{1,0}(P_0), = \pi_{1,1}(P_0)$ and therefore $\pi'_{1,0}(P_0) =$

$\pi_{1,1}'(P_0)$. Rearranging the latter:

$$\begin{aligned}0 &= g_0(P_0)\left[\Pr(\sigma_0)(2f_0-1+P_0)-\Pr(\sigma_1)(2f_1-1+P_0)\right]+ \\ &+\Pr(\sigma_0)\left[G_0(P_0)-G_0(P_0+2\Delta f)\right]+\Pr(\sigma_1)\left[G_0(P_0-2\Delta f)-G_0(P_0)\right]+ \\ &-(2f_1-1+P_0)\Pr(\sigma_0)\,g_0(P_0+2\Delta f)+(2f_0-1+P_0)\Pr(\sigma_1)\,g_0(P_0-2\Delta f)\end{aligned} \tag{7.1}$$

All terms on the r.h.s. are either non positive or strictly negative, except possibly the last one, so we require $g_0(P_0-2\Delta f)>0$, or $P_0-2\Delta f=IC_0(P_1'')\in S_0$. Then $P_1''\in S_1$, $IC_0(P_1'')\in S_0$ and $IC_1(P_1'')=P_0\in S_0$, and the situation that we are trying to rule out for firm 0 at P_0 is replicated for firm 1 at P_1''. By symmetry, this implies $g_1(P_1''-2\Delta f)>0$. (We have verified this implication step by step). Then, recursively $P_0-2\Delta f$ must be in S_0, and it ties with $P_1''=IC_1(P_0-2\Delta f)\in S_1$, as just seen, and $P_1''-2\Delta f=IC_0(P_0-2\Delta f)\in S_1$ by $g_1(P_1''-2\Delta f)>0$. By reiterating the reasoning, initially applied to P_0, this time to $P_0-2\Delta f$, we require $g_0(P_0-4\Delta f)>0$ and, at any further step $n>2$, $g_0(P_0-n\cdot 2\Delta f)>0$, for otherwise the whole argument would unravel. But clearly for $n=N$ large enough and for $(q,\alpha)\in(1/2,1)^2$, so that $\Delta f=f_1-f_0>0$, we must have $P_0-2N\Delta f<0$ and thus $g_0(P_0-2N\Delta f)=0$, giving the desired contradiction.♦

For firm 1, which is favored by the prior belief, we can say even more: given any randomization by firm 0, firm 1's profits are increasing faster on IC_0 than on IC_1:

Lemma 6 (From up to down). *For almost all $P_0\leq 1-f_1$ in the interior of S_0 the profit of seller 1 increases less along IC_1 than along IC_0 in the price of the competitor:*

$$\frac{d\pi_1(IC_1(P_0))}{dP_0}<\frac{d\pi_1(IC_0(P_0))}{dP_0}.$$

Proof. To establish the claim it suffices to prove that $g_0(P_0-2\Delta f)=0$ and use equation (7.1). By Lemma 5, ignoring zero Lebesgue measure sets, P_1' and P_1'' cannot be both in S_1. However at least (and therefore *exactly*) one of the two is, otherwise firm 0 would strictly gain by deviating from P_0 to some $P_0+\varepsilon$. Suppose by contradiction that $g_0(P_0-2\Delta f)>0$. If $P_1'\in S_1$ and thus $P_1''\notin S_1$, we contradict Lemma 5, as $P_0=IC_0(P_1')\in S_1$ and $P_0-2\Delta f=IC_1(P_1')\in S_1$. If instead $P_1''\in S_1$ and thus $P_1'\notin S_1$, as $P_1'=IC_1(P_0-2\Delta f)\notin S_1$ while $P_0-2\Delta f\in S_0$, as again one of the two projections *must* be in the opponent's support, it follows that $P_1'''=IC_0(P_0-2\Delta f)\in S_1$. But $P_1'''=2f_0-1+P_0-2\Delta f=P_1'-2\Delta f<P_1'$, so we have P_1''', $P_1''\in S_1$, and $P_1'\notin S_1$, where $P_1'\in(P_1''',P_1'')$. Since we are concerned only with non zero Lebesgue measure sets of such prices, this requires $(P_1'-\delta,P_1'+\delta)\cap S_1=\emptyset$ for some $\delta>0$. No hole in S_0 may correspond through IC_0 to the hole $(P_1'-\delta,P_1'+\delta)$ in S_1, because the former would contain $(P_0-\delta,P_0+\delta)$ while $(P_0-\varepsilon,P_0+\varepsilon)\subset S_0$ by assumption; so by Lemma 4 a hole in S_0 must project $(P_1'-\delta,P_1'+\delta)$ through the *other* constraint IC_0, and be of the form $(P_0 - \delta - 2\Delta f, P_0 + \delta + 2\Delta f)$. This hole contains $P_0 - 2\Delta f$, contradicting $g_0(P_0 - 2\Delta f) > 0$.♦

Proof of Proposition 4 (Uniqueness of Equilibrium).

In the pooling region there are a continuum of non-cautious equilibria, while the cautious pooling equilibrium is always unique. The separating equilibrium is easily seen to be the unique equilibrium in the separating region. The various mixed-strategy equilibria constructed are mutually exclusive by construction, other than possibly at the boundaries between the different regions. To establish uniqueness we need to exclude mixed strategy equilibria that do not fall into one of the four classes M1-4.

From claims (i) and (ii) of Lemma 3 in equilibrium there can be at most two atoms at the maximum prices of the relevant range and two corresponding gaps, and the rest of the support is connected. Therefore, using Lemma 5 and Lemma 6, M4 is the only possible equilibrium if there are no atoms.

If there are both atoms the equilibrium is of type M1 or M2, since the other a priori possible forms of the equilibrium (similar to M1 with the indexes of the sellers interchanged) do not exist for $q \geq 1/2$, as we now show by contradiction. By symmetry with respect to M1 (details missing, available on request), such an equilibrium would have $\underline{P}_0 = \alpha(1-q)$, $\underline{P}_1 = 2f_1 - 1 + \alpha(1-q)$, $V_0 = \underline{P}_0$, and $V_1 = \Pr(\sigma_1)\underline{P}_1$, with $\underline{P}_0 < 1 - f_1$ by construction. The atom on $\overline{P}_0 = 1 - f_0$ would have measure $\gamma_0 = \underline{P}_1 / f_1$, so that in order for 1 not to deviate to f_0, one needs $\pi_1(f_0) = f_0(\Pr(\sigma_1) + \gamma_0 \Pr(\sigma_0)) < V_1$, equivalent to $\underline{P}_1 > f_0/(2\alpha - 1)$. But this inequality is incompatible with $\underline{P}_0 < 1 - f_1$ (equivalent to $q < 1 - q^S(\alpha)$) for $q \geq 1/2$.

Finally, from Lemma 5 again, M3 is the only possible form of the equilibrium which features only an atom by firm 0 on $P_0 = 1 - f_0$. So we are left to exclude the symmetric equilibrium with only an atom by firm 1 on $\overline{P}_1 = f_1$, which is the hardest case. By contradiction. By Lemma 2, in such an equilibrium firm 0 must play with positive probability prices above $1 - f_1$. By Lemma 4, to this (atomless) mass there must correspond a mass by seller 1 below f_0 . If $2f_1 - 1 > f_0$, firm 1 randomizes below f_0 and on $\overline{P}_1 = f_1$, but not in the interval $[f_0, 2f_1 - 1]$ dominated by $2f_1 - 1$, so that there would be a hole in S_1 without an atom by firm 0, contradicting Lemma 3. If instead $2f_1 - 1 \leq f_0$, absent atoms by firm 0, firm 1 must play below and above f_0, with no hole in S_1. Firm 0 must play prices below $1 - f_1$, and then below the hole, for otherwise prices in (f_0, f_1) would be dominated by f_1 for firm 1, and there would be a hole in S_1. Firm 0 must play all the way down to $1 - 2f_1 + f_0$ to avoid this hole, and not below it, otherwise Lemma 6 would apply. Therefore the following four prices are in the support of firm 0: $1 - 2f_1 + f_0$, $1 - f_1$, $\tilde{P}_0$, and, $1 - f_0$. Equating the four payoffs yields a system of three equations in the three unknowns: $\tilde{P}_0$, the measure of firm 1's atom, and the probability mass played by firm 1 in $[f_0, f_1)$, which we have shown must be positive. Given the solution for $\tilde{P}_0$, consider prices in S_1: $2f_0 - 1 + \tilde{P}_0$, f_0, f_1. Equating the expressions for the payoffs of firm 1 at these prices yields two independent equations to determine one unknown only, the fraction of the probability played by firm 0 above the hole. Therefore the system is overdetermined and has no solution.♦

Proof of Proposition 8 (Degenerate menus)

As mentioned in the discussion above, it is enough to show that the best reply of a seller to any menu of lotteries of the opponent is a degenerate menu. Consider first the case where no type is excluded, i.e. $\lambda_0, \lambda_1 \neq 0$. Adapting standard arguments from non-linear pricing with the twist of type dependent reservation utilities (with $u_0 \geq u_1$), it can be easily shown that PC_0 and SS_1 are always binding. In detail: (1) PC_0 and SS_1 imply PC_1. The proof of this is standard, other when making use of $u_0 \geq u_1$. (2) PC_1 is slack since $\lambda_0 > 0$. (3) PC_0 must be binding, for otherwise T_0 and T_1 could be reduced by the same small amount without affecting the other constraints but increasing profits. (4) SS_1 is binding, for if both PC_1 and SS_1 were slack T_1 could be reduced by a small amount without violating the other constraints but increasing profits.

Substituting $T_0 = \lambda_0 f_0 - u_0$ and $T_1 = (\lambda_1 - \lambda_0) f_1 + T_0$, the objective function becomes

$$\max_{\lambda_0, \lambda_1 \in [0,1]} \Pr(\sigma_1)\lambda_1 + (f_0 - \Pr(\sigma_1) f_1)\lambda_0 - u_0.$$

For it to be optimal not to exclude the type-0 buyer ($\lambda_0 > 0$) it is necessary that $q \geq q^S$ (defined in Section 4.2.1). The optimal non-excluding menu $\{\lambda_0 = 1, T_0 = f_0 - u_0\}$, $\{\lambda_1 = 1, T_1 = f_0 - u_0\}$ is degenerate and corresponds to an unconditional pooling price offer. It yields $\pi = f_0 - u_0$. The remaining constraint SS_0 is immediately checked to be satisfied.

Conditional on excluding the type-0 buyer by setting $\lambda_0 = 0$, PC_1 — rather than SS_1 — is binding, and the optimal lottery is degenerate with $\{\lambda_1 = 1, T_1 = f_1 - u_1\}$, corresponding to the separating price, and $\pi = \Pr(\sigma_1)(f_1 - u_1)$. Comparing profits it is immediately seen that the separating menu is preferred to the pooling menu if and only if $\Pr(\sigma_1 | q)[f_1(q, \alpha) - u_1] \geq f_0(q, \alpha) - u_0$.♦

A.II. Construction of Equilibrium Type M1

Description. The equilibrium prices that define the support, payoffs and randomizations of M1 are $\{\tilde{P}_j, \overline{P}_j, \underline{P}_j\}$, $j = 0, 1$, which solve *uniquely* the system:

$$\underline{P}_1 = \alpha q, \quad \tilde{P}_1 = f_0, \quad \overline{P}_1 = f_1, \quad \underline{P}_0 = \tilde{P}_0 = 1 - 2f_1 + \underline{P}_1, \quad \overline{P}_0 = 1 - f_0.$$

Seller 0 randomizes over an interval of prices $[\tilde{P}_0, 1 - f_0)$, where $\tilde{P}_0 = \underline{P}_0 > 1 - f_1$, with an atomless distribution G_0 of total mass $G_0(f_0) = 1 - \gamma_0$, and on the highest possible price

$1 - f_0$ with an atom of probability mass

$$\gamma_0 = \frac{\alpha}{1-\alpha} - \frac{\Pr(\sigma_1)}{\Pr(\sigma_0)}; \tag{7.2}$$

seller 1 randomizes on $[\underline{P}_1, f_0]$ with an atomless distribution G_1 of total mass $G_1(f_1) = 1 - \gamma_1$, and on f_1 with an atom of probability mass

$$\gamma_1 = \frac{1 - 2f_0 + \alpha q}{1 - f_0}. \tag{7.3}$$

According to the chosen tie-breaking rule, at prices $\{1 - f_0, f_0\}$ the indifferent type-0 buyer chooses good 1 since firm 0 posts price f_0 with positive probability; at prices $\{1 - f_1, f_1\}$ the (indifferent) type-1 buyer to the low price seller 0. The equilibrium payoffs to the sellers are:

$$V_0(q,\alpha) = \Pr(\sigma_0|q,\alpha)\left[1 - 2f_0(q,\alpha) + \alpha q\right], \quad V_1(q,\alpha) = \alpha q. \tag{7.4}$$

To verify that these strategies and tie-breaking rule constitute an equilibrium, we show that all prices in the support of the probability distribution of each seller yield the same expected payoff, given the strategy of the other seller, and that no seller has a strictly profitable deviation to prices not in the support.

Payoffs. The expected payoffs associated with three benchmark prices in the support of firm 0 (firm 1) given the seller 1's strategy G_1 stated above (resp. firm 0's strategy G_0) are easily computed with the help of Figure 6:

Price P_0	Payoff $\pi_0(P_0)$	Price P_1	Payoff $\pi_1(P_1)$
$\underline{P}_0 = \widetilde{P}_0$	$\Pr(\sigma_0)\widetilde{P}_0$	$\underline{P}_1$	$\underline{P}_1$
$1 - f_0$	$\gamma_1 \Pr(\sigma_0)(1 - f_0)$	f_0	$[(1-\gamma_0)\Pr(\sigma_1) + \gamma_0]f_0$
		f_1	$\Pr(\sigma_1) f_1 = \alpha q.$

Solving for Strategies and Randomizations. All prices in the support must yield the same expected payoff V_1: from $\pi_1(\underline{P}_1) = \pi_1(f_1)$ we find $V_1 = \underline{P}_1 = \alpha q$ and from $\pi_1(f_0) = \pi_1(f_1)$ we obtain (7.2), the mass of the atom γ_0 played by seller 0 on the maximum price $1 - f_0$. Notice that $d\gamma_0/dq < 0$, $\gamma_0(q^S) = 1$, and $\gamma_0(1) = 0$. The prices $P_1 \in (\underline{P}_1, f_0)$ in the support are left to be considered. The randomization G_0 of seller 0 must be such that seller 1 is indifferent among all such prices in the support of G_1 which yield: $V_1 = [1 - G_0(1 - 2f_0 + P_1) + G_0(1 - 2f_0 + P_1)\Pr(\sigma_1)]P_1$. Equating this to $V_1 = \alpha q$ and substituting $P_0 = 1 - 2f_0 + P_1$, we obtain $G_0(P_0) = (P_0 + 2f_0 - 1 - \alpha q) / \{\Pr(\sigma_0)[2f_0 - 1 + P_0]\}$. Notice that $G_0(\underline{P}_0) = 0$ and that the density played by seller 0, $g_0(P_0) = G'(P_0)$, is strictly decreasing in P_0.

Substituting back $\tilde{P}_0 = 1 - 2f_0 + \underline{P}_1 = 1 - 2f_0 + \alpha q$ in the two equations above we obtain the equilibrium payoff for firm 0 given in (7.4) and the mass of the atom in (7.3). Notice that $\gamma_1(q) < 1$ for $q \in [q^S, q^P]$. The c.d.f. G_1 played by seller 1 must make seller 0 indifferent among all the remaining prices in the stated support $P_0 \in (\underline{P}_0, 1 - f_0)$ which yield $V_0 = P_0 \Pr(\sigma_0)[1 - G_1(2f_0 - 1 + P_0)]$. Equating this to (7.4) obtained above and substituting $P_1 = 2f_0 - 1 + P_0$, we obtain $G_1(P_1) = (P_1 - \alpha q)/(1 - 2f_0 + P_1)$. Notice that $G_1(\underline{P}_1) = 0$ and $G_1(f_0) = 1 - \gamma_1$, so that $G_1(f_1) = 1$. The density $g_1(P_1) = G_1'(P_1)$ is decreasing in P_1.

Deviations. Given seller 0's strategy, any price less than $\underline{P}_1$ is dominated by $\underline{P}_1$, and any $P_1 \in (f_0, f_1)$ by f_1, as immediately seen from Figure 6. Given seller 1's strategy, seller 0's best deviation is $P_0' = 1 - f_1$: any price below P_0' would result in the same probability of selling as P_0' but at a lower price, and similarly for a price between P_0' and $\tilde{P}_0$ compared with $\tilde{P}_0$. This best deviation is not profitable provided that $\pi_0(P_0') = [\gamma_1 + (1 - \gamma_1)\Pr(\sigma_0)](1 - f_1) \leq V_0$. By (7.4), this is equivalent to $\psi(q, \alpha) := (2\alpha - 1)(1 - 2f_0 + \alpha q) - \alpha(1 - f_1) \geq 0$. First, notice that $\psi''(q) = -2(2\alpha - 1)f_0''(q) + \alpha f_1''(q) < 0$ because $f_0''(q) > 0 > f_1''(q)$ for all q, so that ψ is strictly concave and quasi-concave in q: hence $\psi(q) \geq 0$ for q belonging to an interval Q_{M1}. Next, it can be verified that for $\alpha > 2/3$, i.e. whenever a separating equilibrium exists, $\psi(q) = 0$ has a unique root $q^{M1} > q^S$. For $q < q^S$ the atom on the separating price $P_1 = f_1$ in the M1 equilibrium would have a mass exceeding 1 (cf. 7.3), so that M1 may exist only in $Q_{M1} = [q^S, q^{M1}] \subset [q^S, q^P]$. Instead, for $\alpha < \frac{2}{3}$ there are two roots (q^{m1}, q^{M1}), with $q^{M1} < q^{M1}$ and $\psi'(q^{m1}) < 0 < \psi'(q^{M1})$, so that M1 may exist only in $Q_{M1} = [q^{m1}, q^{M1}] \subset [1/2, q^P]$.

Construction of Mixed Strategy Equilibrium M2-3-4. The same procedure leads to construct the equilibrium strategies in the other three equilibrium regimes, and to show that each pair is indeed an equilibrium only in the corresponding region of parameters, that we depicted in Figure 3 with the help of a Mathematica^{TM} routine.

References

[1] BERGEMANN, DIRK AND JUUSO VALIMAKI, "Learning and Strategic Pricing," *Econometrica*, 1996, 64-5, 1125-1149.

[2] BERGEMANN, DIRK AND JUUSO VALIMAKI, "Market Experimentation and Pricing," 1997, Yale University and Northwestern University mimeo.

[3] BERTRAND, JOSEPH, "Theorie Mathematique de la Richesse Sociale," *Journal de Savants*, 1883, 499-508.

[4] BIKHCHANDANI, SUSHIL, DAVID HIRSHLEIFER AND IVO WELCH, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, 100, 992-1026.

[5] BÖRGERS, TILMAN, "Iterated Elimination of Dominated Strategies in a Bertrand-Edgeworth Model," *Review of Economic Studies*, 1992, 59, 163-176.

[6] BROECKER, THROSTEN, "Credit-Worthiness Tests and Interbank Competition," *Econometrica*, 1990, 58-2, 429-452.

[7] BUTTERS, GERARD, "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies*, 1977, 44, 465-491.

[8] CAMINAL, RAMON AND XAVIER VIVES, "Why Market Shares Matter?: An Information-Based Theory," *RAND Journal of Economics*, 1996, 27-2, 221–239.

[9] CRÉMER, JACQUES, FARAD KHALIL AND JEAN-CHARLES ROCHET, "Contracts and Productive Information Gathering," 1997, GREMAQ Cahier n. 97.25.468.

[10] D'ASPREMONT, CLAUDE, JEAN J. GABSZEWICZ AND JACQUES-FRANCOIS THISSE, "On Hotelling's "Stability in Competition"," *Econometrica*, 1979, 47-5, 1145-1150.

[11] DASGUPTA, PARTHA AND ERIC MASKIN, "The Existence of Equilibria in Discontinuous Economic Games, 1: Theory," *Review of Economic Studies*, 1986, 53, 1-27.

[12] EPSTEIN, LARRY G. AND MICHAEL PETERS, "A Revelation Mechanism for Competing Mechanism," 1998 University of Toronto mimeo.

[13] FELLI, LEONARDO AND CHRISTOPHER HARRIS, "Learning, Wage Dynamics, and Firm-Specific Human Capital," *Journal of Political Economy*, 1996, 104(4), 838-68.

[14] GABSZEWICZ, JEAN J. AND ISABEL GRILO, "Price Competition When Consumers Are Uncertain about Which Firm Sells Which Quality," *Journal of Economics and Management Strategy*, 1992, 4, 629-650.

[15] HARRINGTON, JOSEPH E., "Experimentation and Learning in a Differentiated-Products Duopoly," *Journal of Economic Theory*, 1995, 66, 275-288.

[16] HOTELLING, HAROLD, "Stability in Competition," *Economic Journal*, 1929, 39, 41-57.

[17] MANELLI, ALEJANDRO AND DANIEL VINCENT, "Optimal Procurement Mechanisms," *Econometrica*, 1995, 63-3, 591-620.

[18] MCAFEE, PRESTON R., "Mechanisms Design by Competing Sellers," *Econometrica*, 1993, 61-6, 1281-1312.

[19] MILGROM, PAUL R. AND ROBERT J. WEBER, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, 50-5, 1089-1122.

[20] MOSCARINI, GIUSEPPE AND MARCO OTTAVIANI, "Social Learning and Competition," 1997, Yale University and University College London mimeo.

[21] OSBORNE, MARTIN AND CAROLYN PITCHIK, "Equilibrium in Hotelling's Model of Spacial Competition," *Econometrica*, 1987, 55-4, 911-922.

[22] OTTAVIANI, MARCO, "Monopoly Pricing with Social Learning," 1997, University College London mimeo.

[23] ROSENTHAL, ROBERT W. AND ANDREW WEISS, "Mixed-Strategy Equilibrium in a Market with Asymmetric Information," *Review of Economic Studies*, 1984, 51, 333-342.

[24] SCHLEE, EDWARD E., "The Value of Information about Product Quality," *RAND Journal of Economics*, 1996, 27-4, 803–815.

[25] SHAKED, AVNER AND JOHN SUTTON, "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, 1982, 49, 3-13.

[26] SIMON, LEO AND WILLIAM ZAME, "Discontinuous Games and Endogenous Sharing Rules," *Econometrica*, 1990, 58, 861-872.

[27] STOLE, LARS, "Nonlinear Pricing and Oligopoly," *Journal of Economics and Management Strategy*, 1995, 4, 529-562.

[28] VARIAN, HAL, "A Model of Sales," *American Economic Review*, 1980, 70-4, 651–659.

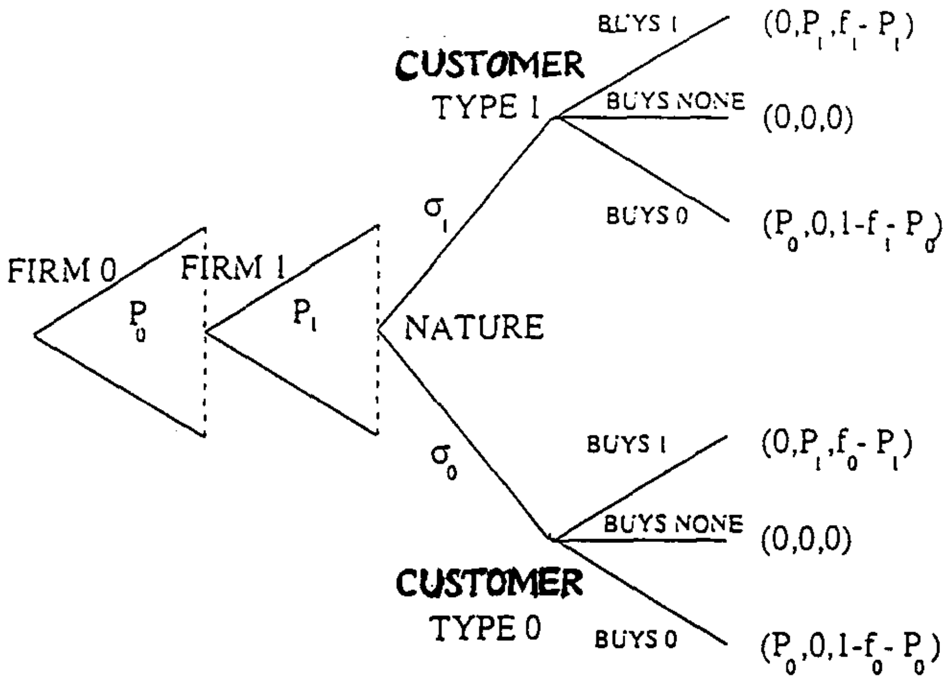

Figure 1: The game.

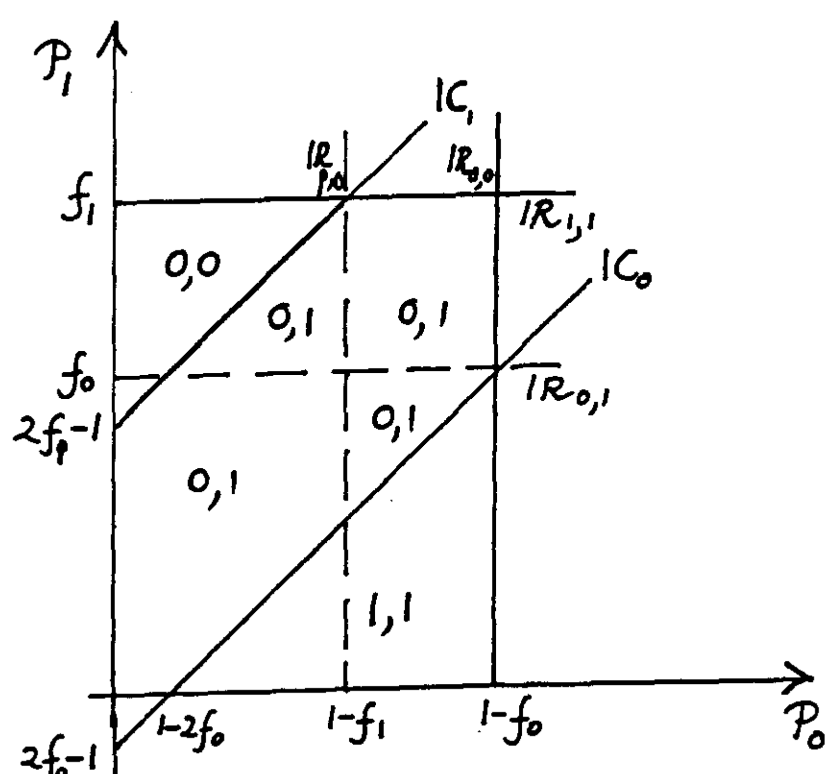

Figure 2a: The constraints in the price space.

Representation on the Hotelling line

FIG. 2b

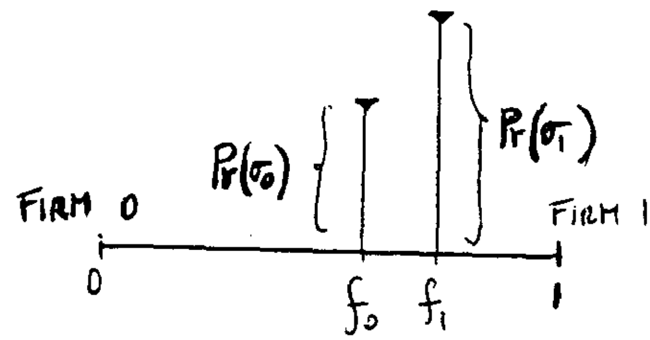

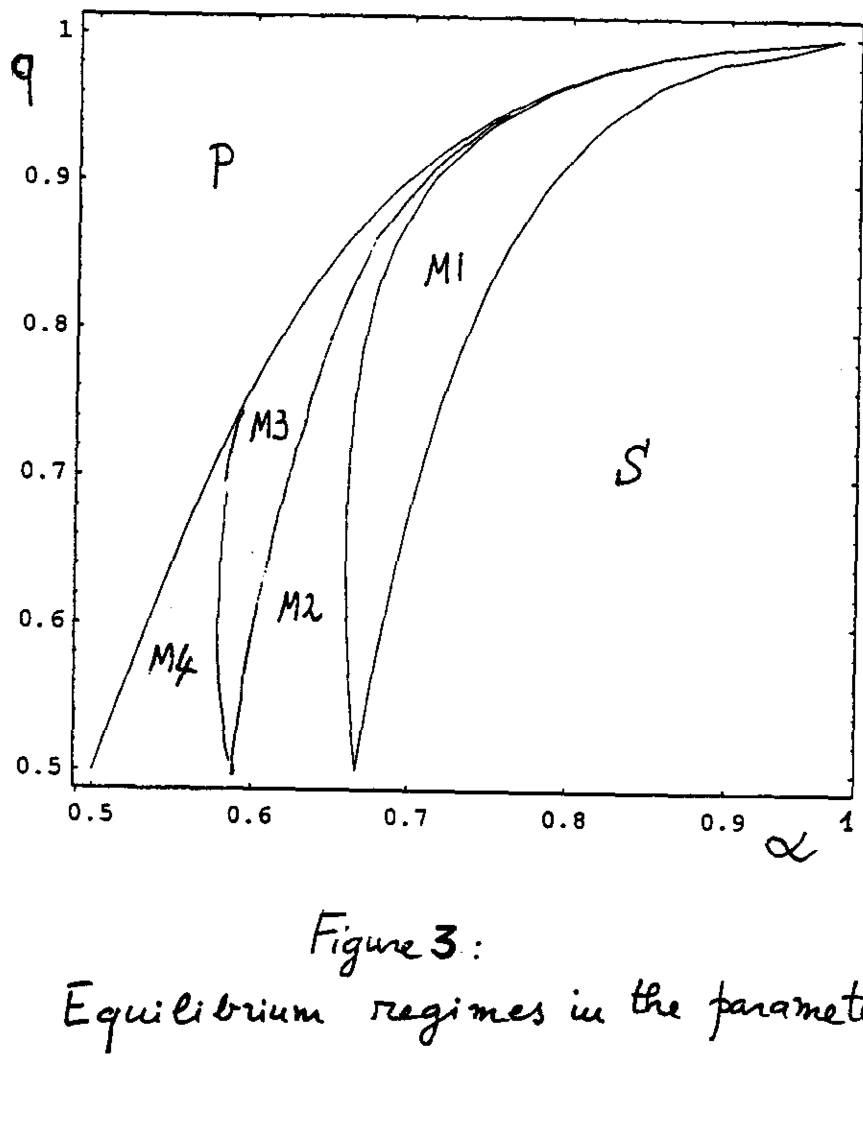

Figure 3:

Equilibrium regimes in the parameter space

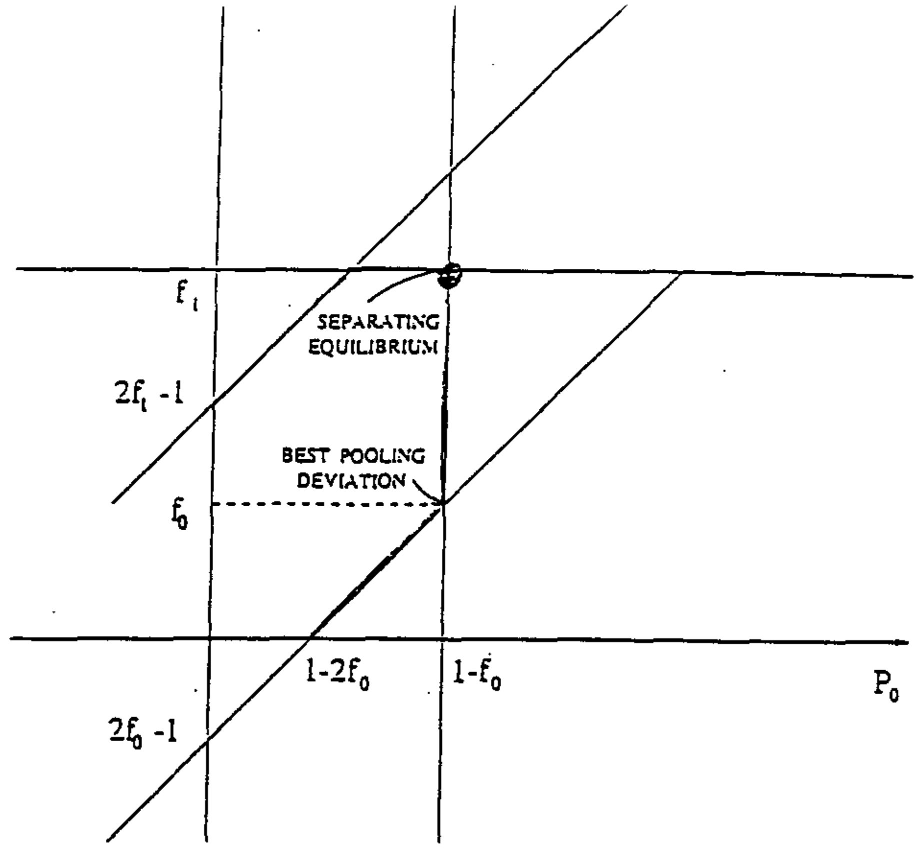

Figure 4:
The separating equilibrium

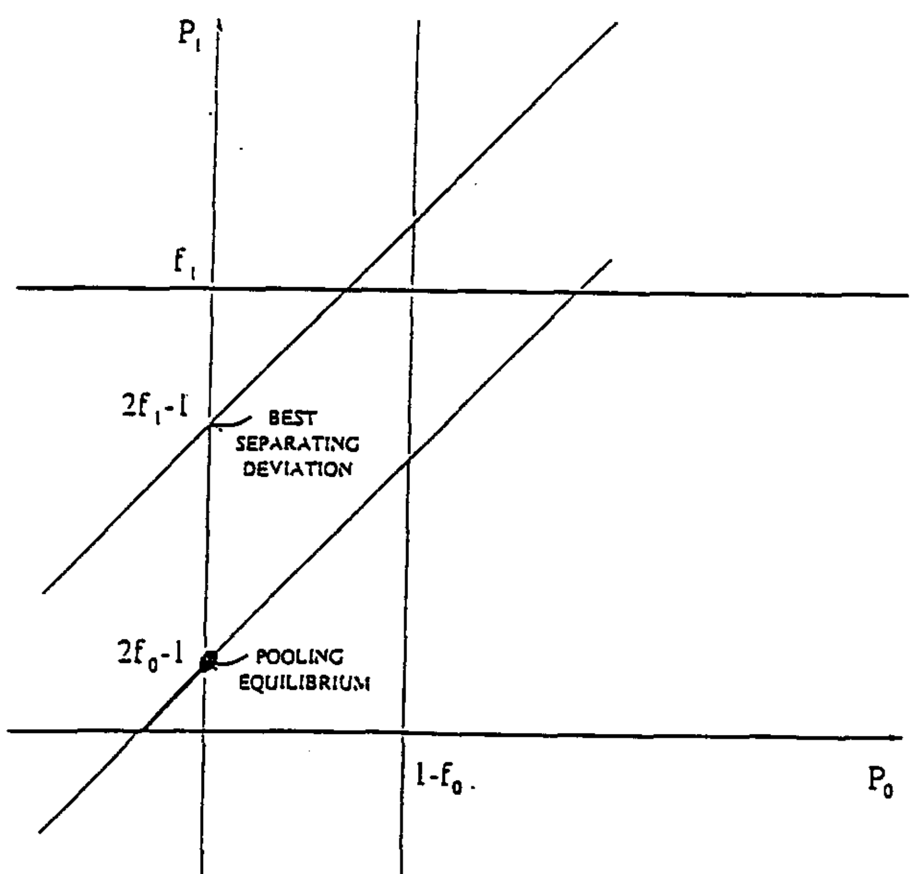

Figure 5:
The pooling equilibrium

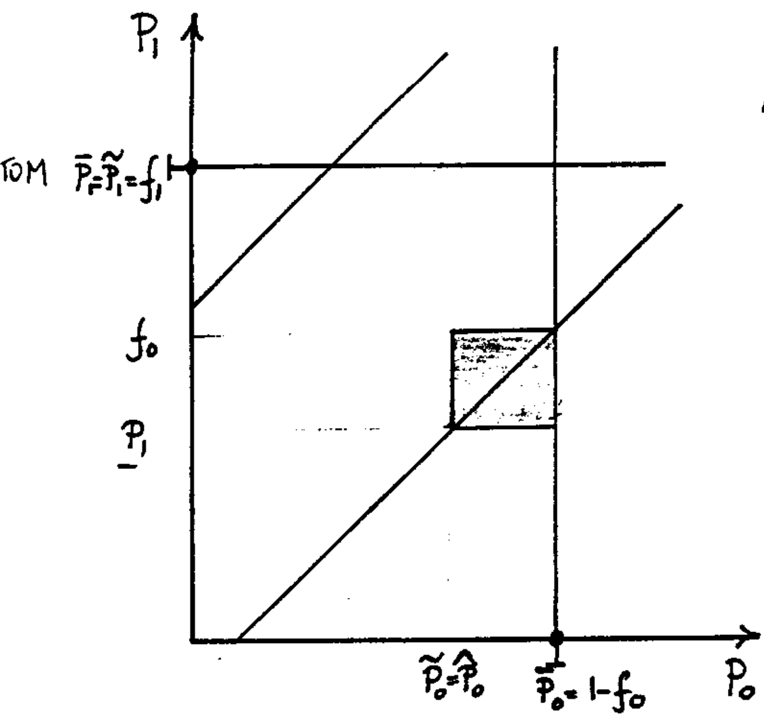

Figure 6 : M1

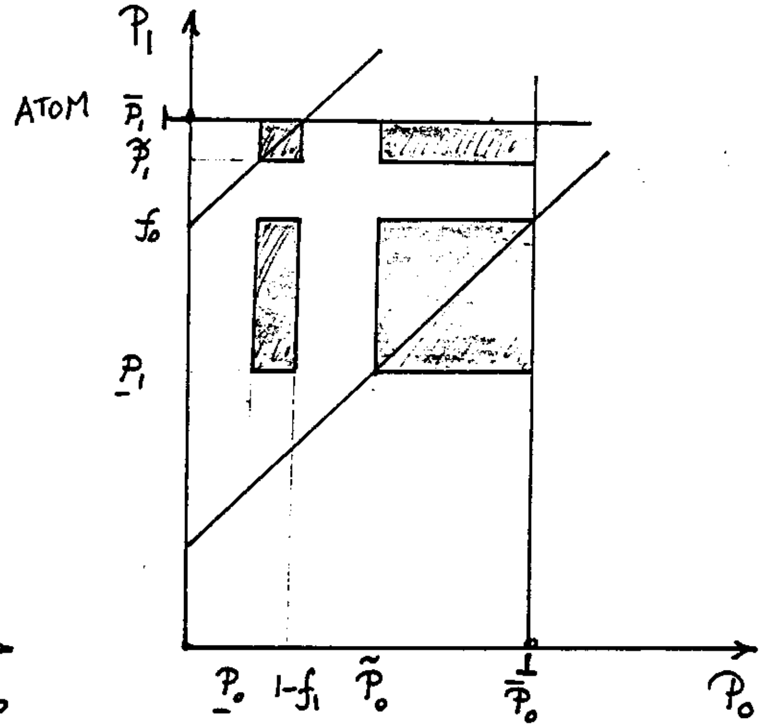

Figure 7: M2

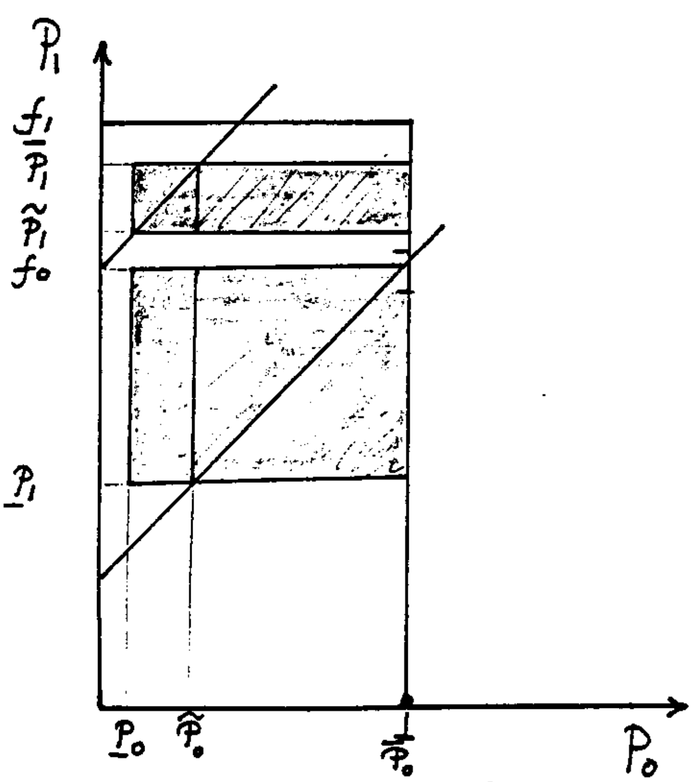

Figure 8: M3

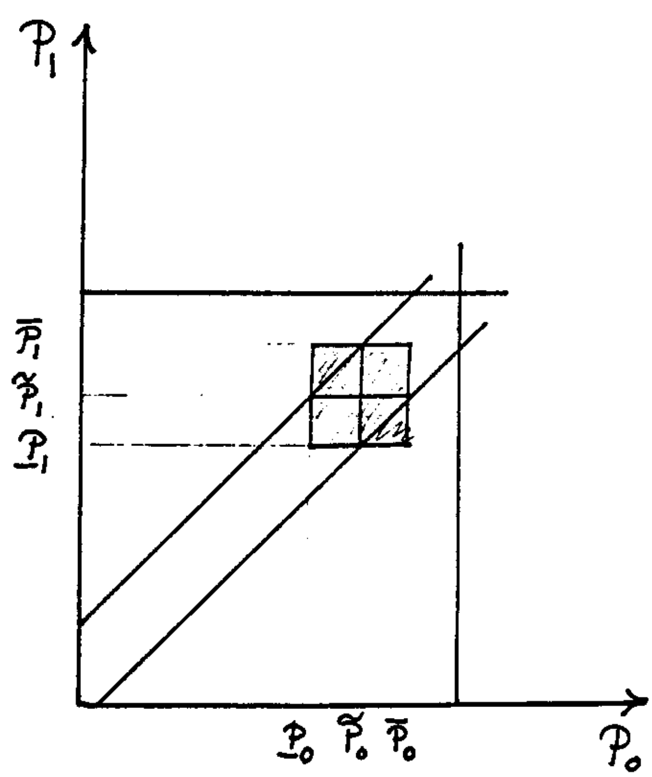

Figure 9 : M4

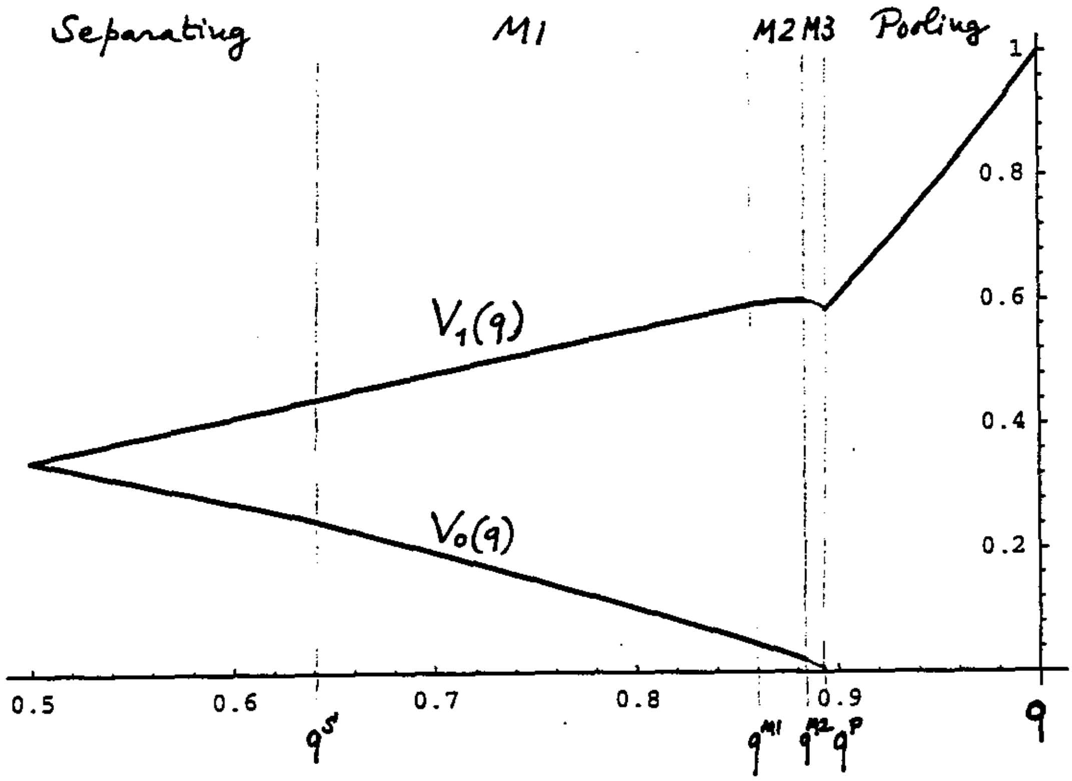

Figure 10. The value functions of the two sellers for $\alpha = .69$ (generated by Mathematica)