On Price versus Quality

Avrim Blum*[1] and Yishay Mansour†[2]

1 Toyota Technological Institute at Chicago, USA
avrim@ttic.edu

2 Tel-Aviv University, and Google Research, Israel
mansour.yishay@gmail.com

Abstract

In this work we propose a model where the value of a buyer for some product (like a slice of pizza) is a combination of their personal desire for the product (how hungry they are for pizza) and the quality of the product (how good the pizza is). Sellers in this setting have a two-dimensional optimization problem of determining both the quality level at which to make their product (how expensive ingredients to use) and the price at which to sell it. We analyze optimal seller strategies as well as analogs of Walrasian equilibria in this setting. A key question we are interested in is: to what extent will the price of a good be a reliable indicator of the good's quality?

One result we show is that indeed in this model, price will be a surprisingly robust signal for quality under optimal seller behavior. In particular, while the specific quality and price that a seller should choose will depend highly on the specific distribution of buyers, for optimal sellers, price and quality will be linearly related, independent of that distribution. We also show that for the case of multiple buyers and sellers, an analog of Walrasian equilibrium exists in this setting, and can be found via a natural tatonnement process. Finally, we analyze markets with a combination of "locals" (who know the quality of each good) and "tourists" (who do not) and analyze under what conditions the market will become a tourist trap, setting quality to zero while keeping prices high.

1998 ACM Subject Classification F.0 General

Keywords and phrases Algorithmic game theory, pricing, Cobb-Douglas valuations

Digital Object Identifier 10.4230/LIPIcs.ITCS.2018.16

1 Introduction

Buyers often use price as a simple proxy for quality: "if it's more expensive, it must be better". This can produce seemingly irrational behavior, such as purchasing a more expensive good or service (like a bottle of wine or a hotel room) even if no other information besides price is known. In fact, filters on hotel reservation web sites such as `hotels.com` even allow one to specify a minimum price in addition to a maximum price for searching hotel rooms, which seems somewhat strange under usual models for buyers. Why would anyone refuse a cheaper price? Would the web site be better-off by simply increasing the hotel rooms' prices to the minimum specified price? This can be viewed as irrational behavior by users

* AB was supported in part by NSF grants CCF-1800317, CCF-1525971 and CCF-1331175. This work was conducted in part while the author was visiting the Simons Institute for the Theory of Computing and in part while the author was at Carnegie Mellon University.

† YM was supported in part by a grant from the Israel Science Foundation (ISF), a grant from the United States-Israel Binational Science Foundation (BSF), and the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11). This work was conducted in part while the author was visiting the Simons Institute for the Theory of Computing.

that might be targeted and explained by behavioral economics. Our goal in this work is to understand, model, and analyze the effect of this phenomenon within the assumptions of rationality.

In this work, we propose a model where buyers have private "desires" for goods (as in the usual auction setting [8]) but their true value also depends on the intrinsic quality q of the good (which might or might not be observable to them). Here, we *define* quality as the cost to the seller to produce that good or service, so a higher-quality good is by definition more expensive to make. Sellers, in our model, have the ability to choose both the price and quality level of the goods they aim to sell. We then analyze how markets of this kind behave. For instance, under what conditions will the price of a good indeed be a reliable signal of its quality, and how sensitive is this signal to the distribution of buyer valuations? What fraction of maximum social welfare can sellers extract in such a model? And do Walrasian equilibria exist and can they be produced by some analog of the usual tatonnement process? More broadly, our aim in this work is to think about (seemingly) irrational behaviors of buyers and to understand how robust various classic economics results are to these deviations from the standard setup for them.

1.1 Our Basic Model

We propose a model where the value that a buyer has for a product (like a slice of pizza) depends on both his personal desire for the product (how hungry he is for pizza) and the quality of the product (how good the pizza is). Our basic model builds on the well known Cobb-Douglas production and valuation function (see [7]). In the Cobb-Douglas function for k quantities, the value is $\prod_{i=1}^{k} x_i^{\alpha_i}$, where $\sum_{i=1}^{k} \alpha_i = 1$. For example, the output of an economy is often modeled as a Cobb-Douglas function over labor and capital, and Cobb-Douglas functions are often used to model utilities over combinations of different goods. We similarly, consider the case that the two quantities are the private valuation and the quality, and specifically focus on the case that the value of a buyer is the geometric mean of the two. Formally, each buyer i has some *intrinsic value* v_{ij} for each product j, and if product j has quality q_j then buyer i's valuation will be $\sqrt{v_{ij}q_j}$. We call this *geometric-mean valuations*, which is a Cobb-Douglas valuation with $\alpha_i = 1/2$. (We also generalize our results to buyers with α-*geometric-mean valuations*, in the spirit of the Cobb-Douglas valuation, which are of the form $v_{ij}^{\alpha} q_j^{1-\alpha}$. Note that for $\alpha = 1$ we get the "standard" case of private valuations.)

We assume that higher quality products are more expensive for the seller to produce. In fact, since it's unclear how to define "units of quality" anyway, we simply *define* quality to be the cost to the seller for making a good of that quality. That is, a slice of pizza that costs twice as much to make is twice the quality by definition.

Finally, we assume quasilinear utilities. So, the *utility* to a buyer of intrinsic value v for a good of quality q that is priced at p is:

$$\sqrt{vq} - p.$$

For example, if a buyer has intrinsic value 10 for a slice of pizza, and a pizzeria is selling slices of quality 2.5 at price 4, then the buyer's utility would be $\sqrt{25} - 4 = 1$ and the seller's profit would be $4 - 2.5 = 1.5$. The overall social welfare would be 2.5, the sum of buyer utility and seller profit.

We consider this model in two classic settings. The first is where there is a single seller and some distribution $\mathcal{F}$ over the nonnegative reals from which buyer "intrinsic values" v are drawn. Assuming that quality *is* observable by buyers, the question then is what price and quality should a seller choose in order to maximize profit, and to what extent will price and quality be related in a clean way? The second setting is a market with multiple unit-demand buyers and multiple single-item sellers. In this case, the natural questions involve the existence of market equilibria along with natural procedures for finding them.

1.2 Rationale for Cobb-Douglas valuations

There are many ways one can combine an intrinsic value and a quality to derive a valuation, and using the Cobb-Douglas valuations is one way to do it. One benefit of using the Cobb-Douglas valuation is its long history and usefulness in Microeconomics. In this section we will motivate why such Cobb-Douglas valuations are a better fit to our proposed setting than other simple combination rules.

Taking the geometric mean of the intrinsic value v and quality q has several appealing properties. First, if you are making the good for your own consumption (buying ingredients to make pizza at home) your utility $\sqrt{vq} - q$ is maximized at $q = v/4$, giving utility equal to $v/4$. So the maximum social welfare is linearly related to v. Second, we can think of v as the maximum price a buyer would be willing to pay if a good were priced at cost (i.e., if the price is set to q then a buyer would purchase only quality $q \leq v$, since otherwise he will have a negative utility). Third, the function is monotone in both v and q.

The alternative of taking a linear combination of v and q is more problematic. If the buyer's valuation for a good of quality q were, say, $(v + q)/2$, then social welfare and seller profit would be optimized at $q = 0$; on the other hand if it were, say, $v + 2q$ then social welfare and seller profit would be optimized at $q = \infty$. Another alternative might be to combine on a logarithmic scale: setting valuation to be $\ln(vq)$. However, in this case, social welfare and seller profit would always be optimized at $q = 1$, and there would be no correlation at all between price and quality.

1.3 Our Results

We begin by considering a single seller (such as a pizzeria) who can produce one good (pizza) and wishes to determine the quality level q at which to make the good and the price p at which to sell it. The aim of the seller is to maximize profit given some distribution $\mathcal{F}$ over buyer intrinsic values v. Buyers are assumed to be able to observe quality so that the seller cannot just set $q = 0$. What we show is that in such a setting, price will be a surprisingly good signal for quality. In particular, while the optimal price p^* and quality q^* will depend heavily on the buyer distribution $\mathcal{F}$, their *ratio* is fixed: for all $\mathcal{F}$ we have $q^* = p^*/2$ where p^* and q^* are the price and quality which maximize the seller's profit.

Note that this gives a very simple guideline to sellers on how to price their products. For example, for restaurants, pricing the items on the menu is a potentially very involved task. A common rule of thumb which is widely used is to price a dish at three times the cost of its ingredients.[1] This simple rule of thumb is used across many different geographies, cuisines, and price levels. In part, our results justify the ability to give such a simple universal pricing rule.

We also consider a setting where sellers can produce goods of multiple different quality levels. We give a simple strategy for sellers and we show that it is a constant-factor approximation to optimal for revenue.

[1] See, e.g., `https://www.forbes.com/sites/priceonomics/2017/04/07/how-much-do-the-ingredients-cost-in-your-favorite-foods/#5be51ac611ed`.

Next we turn to the question of market clearing (Walrasian) equilibria. Consider a market with n_b unit-demand buyers and n_g single-good sellers. In the usual model (buyers with private valuations and no quality levels) there always exists a Walrasian equilibrium: an assignment of prices to items and a matching of buyers to items such that (a) every buyer is purchasing his most-desired item at these prices (it is envy-free), and (b) any unpurchased item has price 0. Moreover (c) this matching will maximize social welfare (it is a maximum-weight matching), and (d) it can be found via a natural tatonnement process.[2] We show here that something very similar exists in our model. We begin by proving an analog of (a,b,c). Namely we show that for the social-welfare maximizing assignment of qualities to items and items to buyers, there exist prices that make this assignment an equilibrium.[3] Next we prove an analog of (d), giving a natural process that achieves such an equilibrium.

Finally, we analyze markets that contain both "tourists" and "locals". In these markets, the locals are assumed to know the qualities of various goods or services, whereas the tourists are using price as a signal for quality. For a single good, we analyze how quality will relate to price under optimal seller behavior in such markets and under what conditions the market will become a "tourist trap", setting quality to zero while keeping prices high.

1.4 Related Work

There are many studies in marketing that show how consumers use price as an indicator for quality, and there are empirically-based models that relate perceived quality, price and value, see [12]. There is an ongoing discussion of how a firm can signal an unobservable quality [6]. There is also evidence that the price of products influences the consumer view of their efficacy [9]. (We note that there are also many other works related to quality and pricing from a marketing perspective.)

There are also theoretical works that model the quality-price relationship. The work of [10] studies a market where some of the consumers are price sensitive and some are quality sensitive, and the main issue is that certain consumers might buy products of high price and low quality. The work of [1] shows that in certain instances firms might use an initial high price for a new product, in order to signal to the consumers that the new product is of potentially high quality.

The Cobb-Douglas valuation function is widely studied in Microeconomics (see, [7]), and it has been also applied in various empirical studies. There are empirical studies that relate quality to utility in various domains, such as hospitals [4], mail service [11], soccer coaches [3] and more. In those empirical studies the goal is to find a (best) fit for the Cobb-Douglas parameters, and deduce whether or not the quality plays an important role in influencing the outcome.

2 When price is a signal for quality

Our first result is that in our model of geometric-mean buyers, when there is a single seller who can produce one good (like pizza) at a single quality level, under quite general conditions, price in this setting will in fact be a reliable proxy for quality. This is perhaps surprising since there is no competition.

[2] For the existence of a Walresian equilibrium see [5], and for the analysis of tatonnement see [2]. See also [8].

[3] We define the social welfare as the sum of the buyers' and sellers' utilities, or equivalently, the total value to buyers minus the total cost to sellers.

▶ **Theorem 1.** *For any distribution $\mathcal{F}$ over intrinsic values v, and geometric-mean buyer valuations, the seller maximizes its expected profit per potential customer by choosing price p^* and quality q^* such that $q^* = p^*/2$.*

Proof. Let $\mathcal{F}$ be the distribution over the intrinsic values v of people for the given good. Suppose that the seller chooses some quality level q and price p. Then the seller's expected profit on a potential customer is:

$$(p-q) \Pr_{v\sim\mathcal{F}} [\sqrt{vq} \geq p]$$

or equivalently,

$$(p-q) \Pr_{v\sim\mathcal{F}} \left[v \geq p^2/q\right].$$

Define $\tau = p^2/q$. The profit-maximizing threshold τ will of course depend on specifics of the distribution $\mathcal{F}$. However, notice that once τ is fixed, the probability that a random buyer will purchase the good is fixed as well. This means that for any given τ, the optimal values of p and q for the seller are just the values that maximize $p - q$ subject to $p^2/q = \tau$. We can solve for this by taking the derivative of $p - q = p - p^2/\tau$ with respect to p and setting to 0 to get $1 - 2p/\tau = 0$ or $p^* = \tau/2$. This in turn implies that $q^* = (p^*)^2/\tau = \tau/4$. This implies that for any value of τ we have $q^* = p^*/2$. ◀

The above analysis means that for any distribution $\mathcal{F}$, even though a seller who wants to maximize profit will set both price and quality depending on specifics of $\mathcal{F}$, we will always have price equal to twice the quality at the profit-maximizing solution. This means it is reasonable for buyers to use price as a signal for quality, even if they don't know anything about the market.

We can also relate the outcome in our model compared to the standard quasi-linear model (i.e., with no quality, and utility is $v - p$).

▶ **Claim 2.** *Consider a single good setting and let $z^* = \arg\max_z z \Pr[v \geq z]$ the monopolist maximum revenue price for standard quasi-linear valuations. In the geometric-mean valuation model the revenue-maximizing price will be $p^* = z^*/2$ and quality $q^* = z^*/4$. The ratio between the optimal revenue in the geometric-mean valuation and the standard quasi linear valuation is $1/4$.*

Proof. In the geometric mean model the expected profit is $(p-q)Pr[v > p^2/q]$. Since we have shown that the optimal quality q^* equals $p^*/2$, where p^* is the optimal price, the optimal price is $p^* = \arg\max_p (p/2)Pr[v > 2p]$, which is equivalent to $\arg\max_p (2p)Pr[v > 2p]$, and hence $p^* = z^*/2$, $q^* = z^*/4$. The revenue in this case is $\max_p (p/2)Pr[v > 2p] = (1/4)\max_z z\Pr[v \geq z]$, which implies that the expected revenue in the geometric-mean valuation is a $1/4$ of the revenue in the standard quasi-linear valuation. ◀

We can also consider α-geometric-mean buyers, where the value of a buyer with intrinsic value v for a good of quality q is $v^\alpha q^{1-\alpha}$, i.e., the Cobb-Douglas valuation with parameter α. In this case, we have a linear relation between price and quality where now the slope is $1 - \alpha$.

▶ **Theorem 3.** *For any distribution $\mathcal{F}$ over intrinsic values v, and buyers with α geometric mean valuations, where $\alpha \in (0, 1)$, the seller maximizes its expected profit per buyer by choosing price p^* and quality q^*, where $q^* = (1 - \alpha)p^*$.*

Proof. We can rewrite $\Pr_{v \sim \mathcal{F}}[v^{\alpha} q^{1-\alpha} \geq p]$ as $\Pr_{v \sim \mathcal{F}}[v \geq \tau]$ for $\tau = \frac{p^{1/\alpha}}{q^{(1-\alpha)/\alpha}}$. Fixing τ, we wish to maximize $p - q$. So, plugging in $q = \frac{p^{1/(1-\alpha)}}{\tau^{\alpha/(1-\alpha)}}$ and setting the derivative with respect to p to 0 gives

$$1 - \left(\frac{1}{1-\alpha}\right)\left(\frac{p}{\tau}\right)^{\frac{\alpha}{1-\alpha}} = 0,$$

which solves to $p^* = (1-\alpha)^{\frac{1-\alpha}{\alpha}} \tau$. This in turn implies $q^* = \frac{p^{*\,1/(1-\alpha)}}{\tau^{\alpha/(1-\alpha)}} = (1-\alpha)^{\frac{1}{\alpha}} \tau$, so $q^* = (1-\alpha)p^*$ as desired. ◀

Note that for $\alpha = 2/3$ we get the famous "factor of three" pricing for restaurants.

2.1 Multiple qualities

Suppose a seller has the ability to produce the good at different quality levels, and in fact can produce it at infinitely many quality levels. Then one strategy the seller can use is to produce goods at *all* quality levels, pricing a good at twice its quality. (A more intuitive interpretation is that buyers can select any quality the buyer wishes and pay the seller twice the quality it selected.)

In that case, a buyer with intrinsic value v will choose the item of quality q maximizing $\sqrt{vq} - 2q$, which solves to $q = v/16$ and $p = v/8$. In this case, the buyer will gain a utility of $v/8$ and the seller will gain a profit of $v/16$.

▶ **Theorem 4.** *Consider a seller which offers goods at all quality levels with price $p = 2q$. For any buyer with geometric mean valuation of intrinsic value v, the utility-maximizing price will be $p = v/8$, quality $q = v/16$ and the seller's revenue $v/16$.*

This is interesting in three respects. First of all, notice that the seller cannot hope to get profit greater than $v/4$ even if the seller knows v, since in that case its optimal strategy is to sell one good of price $v/2$ and quality $v/4$ (this is $\tau = v$ in the previous analysis). So, the seller is within a factor of 4 of the best it could possibly hope for, without requiring any information about the buyers distribution. Secondly, the buyer and seller are splitting the surplus roughly equally, which is interesting. Finally, price is still a signal for quality in this case.

Open Question: Out of all possible pricing functions, is $p(q) = 2q$ optimal for the seller in a minimax sense? More specifically, consider a game where the seller selects a pricing function $p(q)$, the adversary selects an intrinsic value v for the buyer, and then the seller's payoff in the game is the fraction of v that it makes in profit when the buyer chooses quality q and price $p(q)$ to maximize its own utility. I.e., the seller's payoff is $(p(q) - q)/v$ where $q = \arg\max_{q'} \sqrt{vq'} - p(q')$. The above analysis shows that $p(q) = 2q$ guarantees the seller a payoff at least $1/16$ in this game. Additionally, we know the value of this game to the seller is at most $1/4$ since even if it sees v in advance and then gets to best-respond, it cannot make more than $1/4$ of v in profit. It is not hard to show that $p(q) = 2q$ is optimal out of possible *linear* functions, but the open question is whether this is optimal out of all possible pricing functions.

3 Unit Demand buyers

In this section we consider a market with multiple goods and multiple buyers. The buyers are unit-demand, which implies that there is always a singleton (or empty) set which is their best response. Our main goal is to show that for buyers with geometric mean valuations we maintain the "nice" equilibrium properties that exist in the standard Walrasian equilibrium setting.

In Section 3.1 we show that for buyers with geometric mean valuations, there always exist prices and qualities that guarantee (a) that each buyer is allocated a best response set, (b) that unsold items have price 0, and (c) that the social welfare is maximized. (Note that for geometric mean valuations, satisfying only (a) and (b) is trivial because we can can set all the prices and qualities to zero, and not allocate any good to any buyer.)

The proof in Section 3.1 does not provide a natural dynamics. In Section 3.2 we improve on this by defining a rather natural two stage dynamics, which we show reaches the desired equilibrium. In this dynamics the buyers compete for the right to receive a good and set its quality.

3.1 Market clearing prices and qualities

Consider n_b unit-demand buyers with geometric-mean valuations and n_g goods. Formally, each buyer i has an intrinsic value $v_{i,j}$ for good j. Given prices p and qualities q, the utility of buyer i for a subset S of goods is,

$$u_i(S) = \max_{j \in S} \sqrt{v_{i,j} q_j} - \sum_{j \in S} p_j.$$

Note that buyer i's utility is always maximized on a set S_i of size at most one.

We would like to consider market clearing prices and qualities. That is, prices and qualities for which each buyer receives a utility-maximizing set S_i, unsold goods have price and quality zero, and overall social welfare is maximized over possible allocations $\{S_i\}$ and qualities $\{q_j\}$. Formally, for each item j either $p_j = q_j = 0$ or $j \in S_i$ for exactly one buyer i, and moreover the social welfare of the allocation is as high as possible where social welfare is defined as

$$\sum_{i=1}^{n_b} u_i(S_i) - \sum_{j=1}^{n_g} q_j = \sum_{i=1}^{n_g} \max_{j \in S_i} \sqrt{v_{i,j} q_j} - \sum_{j=1}^{n_g} q_j \ ,$$

and the maximization is both over the allocations S_i and the qualities q_j.

One useful fact to note is that a necessary condition for maximizing social welfare is that if $S_i = \{j\}$ then we must have $q_j = v_{i,j}/4$. This implies that the maximum social welfare is well defined as a function of the intrinsic values.

▶ **Theorem 5.** *For any set of n_b unit-demand buyers, with geometric-mean valuations, there are prices p and qualities q which clear the market with an allocation that maximizes social welfare.*

Proof. Consider any matching between buyers and goods. If buyer i is matched to good j, then social welfare is maximized at $q_j = v_{i,j}/4$, and the contribution to social welfare of this pair is $v_{i,j}/4$.

This implies that for any matching, the maximum welfare of that matching is exactly the weight of the matching divided by 4. Therefore, to solve for the maximum welfare solution, we can find the maximum weighted matching and then set qualities equal to the buyer's value divided by 4.

Now, we need to define prices that make this an equilibrium. We will show a reduction to a standard unit-demand valuations and solve for prices in that market, i.e., a Walresian

Algorithm 1: Two phase dynamics

phase 1: Tatonnement process:
  Each buyer competes for goods (where the valuation of buyer i for good j is $v_{i,j}/4$)
outcome: each good j is allocated to at most one buyer $b(j)$ with a price p_j.
phase 2: Setting the qualities:
  For each good j, buyer $b(j)$ sets the quality $q_j = v_{i,j}/4$, pays $p_j + q_j = p_j + v_{i,j}/4$ and receives good j with quality q_j.

equilibrium. Consider the maximum matching between buyers and goods, ignoring qualities. We will define for each good j a quality q_j. For any unallocated item we set the quality to 0 (which will imply later price 0). For good j, which the maximum matching allocated to buyer i we set $q_j = v_{i,j}/4$. After we fix the qualities, we define n_b standard unit demand buyers, where the valuation of buyer k for good j is $\widehat{v}_{k,j} = \sqrt{v_{k,j}q_j}$. We now have a market with unit demand buyers, and therefore there exists a Walresian equilibrium with price $\widehat{p}_j$ for each good j (see, [5]).

Therefore, in our unit demand buyers with geometric mean valuation, the market clearing prices are $\langle \widehat{p}_1, \ldots, \widehat{p}_{n_g} \rangle$ and qualities $\langle q_1, \ldots, q_{n_g} \rangle$ ◀

The above proof does not give a "natural" process to determine the prices, but shows that market clearing prices and qualities always exist for unit demand buyers. Below we modify the above construction to produce a more natural tatonnement-like process.

3.2 A natural dynamics process

We consider the following process to set prices and qualities for each good. In this process, buyers are bidding for the "right to control" each good. That is, prices on each good begin at zero and are raised as in the usual tatonnement process, where the quality of each good will later be *determined and paid for by the winning buyer for it.* In particular, if a buyer wins a good j for a price p_j, she can then afterwards set the quality of the good to any desired q_j and pay a total of $p_j + q_j$.

One way to think of this is like an auction by oil companies for drilling rights on land. Once a company i pays p_j for the rights to drill on land j, it then can decide the amount q_j that it will invest to drill, and it will then receive revenue of $\sqrt{v_{i,j}q_j}$, having paid a total of $p_j + q_j$.

Note that once buyer i purchases (the right to control) good j it will then set a quality $q_j = v_{i,j}/4$ in order to maximize its own utility, which will be $\sqrt{v_{i,j}q_j} - q_j - p_j = v_{i,j}/4 - p_j$. Therefore, when buyer i is considering whether it prefers to pay p_j for item j or $p_{j'}$ for item j', it is comparing $v_{i,j}/4 - p_j$ to $v_{i,j'}/4 - p_{j'}$ just as in the usual Walrasian market setting except the values v have been divided by 4.

So, by allowing buyers to bid on the "right to control" each good, we now have a Walrasian market, where the fundamental valuation of buyer i for good j is $v_{i,j}/4$. This will result in the same outcome as having valuations $v_{i,j}$ (with just the clearing prices a factor of 4 lower). This means it is an equilibrium and social welfare maximizing allocation in our setting as well.

Thus we have the following theorem.

▶ **Theorem 6.** *For unit-demand buyers with geometric mean valuation, the two-phase dynamics (of Algorithm 1) converges to market clearing prices which maximize the social* *welfare.*

4 Two populations: Local and Tourists

In this section we go back to the setting of a single good considered in Section 2, but now assume that the populations of buyers is composed of two sub-populations. One sub-population, which we call *locals*, are aware of the quality of the good. The other sub-population, which we call *tourists*, are unaware of the quality of the good. The tourist sub-population uses the price as a proxy for quality, and uses the naive assumption that the quality is half the price (which would be the case under optimal sellers in a market of only locals). We assume the seller is aware of the fraction of tourists and locals, but has to post a single price for both populations. Our goal is to investigate the effect of tourists on the quality and prices that the seller chooses. For example, we would expect to observe effects similar to tourist traps, where the quality is low and the price is high.

Formally, assume that the population is λ fraction locals and $1-\lambda$ fraction tourists. The tourists have a "one time" experience. They assume that the quality of the good is $q = p/2$. (We assume that the behavior of the tourists is non-strategic in this aspect. We leave for future research considering the case that the tourists are aware of λ , the fraction of locals, and consider an equilibrium of price and quality. The main challenge would be to have a model where the price will not completely reveal the quality.)

We will start with the two extreme cases, $\lambda = 1$ and $\lambda = 0$. The first case is *no tourists*, i.e., $\lambda = 1$. We saw that in this case we have $q = p/2$ and the seller selects a price p that maximizes

$$(p-q)\Pr[\sqrt{vq} \geq p] = (p/2)\Pr[v \geq 2p].$$

In the second case we have *only tourists*, i.e., $\lambda = 0$, then we have the seller selecting a price that maximizes

$$p\Pr[\sqrt{vp/2} \geq p] = p\Pr[v \geq 2p].$$

Clearly in both cases we have the same optimal price p^*, but a different seller quality and revenue. In the no tourists case the quality is $p^*/2$ and the revenue per good is $p^*/2$, while in the only tourists case the quality is 0 and the revenue per good is p^*.

4.1 Price at intermediate values of λ

We saw above that for any $\mathcal{F}$, the optimal price at $\lambda = 1$ equals the optimal price at $\lambda = 0$. This brings up the natural question: is that also true for intermediate values of λ?

The answer to this question is "not necessarily". In particular, consider a distribution $\mathcal{F}$ on intrinsic values v that is uniform over $[0, 1]$. In the case of $\lambda \in \{0, 1\}$, profit is optimized when we maximize $p\Pr[v \geq 2p] = p(1-2p)$, which occurs at $p = 1/4$ (with $q = 1/8$ when $\lambda = 1$ and $q = 0$ when $\lambda = 0$). However, at $\lambda = 0.5$, a calculation shows that profit is maximized at $p = 1/(3+\sqrt{3}) \approx 1/4.7$, and $q = p/(1+\sqrt{3})$. For example, with $\lambda = 0.5$, if the seller sets $p = 1/4$ then the maximum profit she can make is $1/16 = 0.0625$ (which occurs both at $q = 0$ and $q = 1/8$). But at the optimal p and q, she gets a (slightly) higher profit of 0.067.

Another natural qualitative question is: what happens to quality when the fraction of locals is small? When the fraction of locals is zero, then clearly the seller's optimal quality is zero. The question is whether or not any strictly positive fraction of locals is sufficient to drive the quality away from zero.

In the next section we show that the answer to this question is distribution-dependent. For many distributions the quality will remain zero, while for some distributions the quality will increase from zero.

4.2 Pareto distribution for intrinsic values

We consider a Pareto distribution for the intrinsic values. Recall that a Pareto distribution has two parameters, $\beta > 0$ and $x_{\min} > 0$. We will set $x_{\min} = 1$, so we have a single parameter. The density of the distribution is $f(x) = \beta/x^{1+\beta}$ and $\Pr[v \geq x] = (1/x)^\beta$.

We can now write the revenue R as a function of the price p, quality q and parameter β, i.e.,

$$R(p, q, \beta) = \lambda(p - q)(\frac{q}{p^2})^\beta + (1 - \lambda)(p - q)(\frac{1}{2p})^\beta$$

Clearly maximizing the revenue implies that $q \in [0, p]$. Consider the derivative of the revenue with respect to the quality,

$$\frac{\partial}{\partial q} R = -\lambda(\frac{q}{p^2})^\beta + \lambda(p - q)\frac{\beta}{p^2}(\frac{q}{p^2})^{\beta-1} - (1 - \lambda)(\frac{1}{2p})^\beta$$

For $\beta = 1$ we have

$$\frac{\partial}{\partial q} R = \frac{\lambda}{p} - \frac{2\lambda q}{p^2} - \frac{1 - \lambda}{2p}$$

We get that the optimal q^* is

$$q^* = \max\{0, \frac{p(3\lambda - 1)}{4\lambda}\}$$

which for $\lambda \leq 1/3$ implies that $q^* \leq 0$ and hence the optimal q for $\lambda \leq 1/3$ is $q^* = 0$.
For $\beta = 1/2$ we have

$$\frac{\partial}{\partial q} R = \frac{\lambda}{2\sqrt{q}} - \frac{3\lambda\sqrt{q}}{2p} - \frac{1 - \lambda}{\sqrt{2p}}$$

In this case we show that even for $\lambda \approx 0$ we will have $q^* > 0$. The optimal quality q^* is

$$\begin{aligned}
\sqrt{q^*} &= \frac{-\sqrt{2}(1 - \lambda)/\sqrt{p} \pm \sqrt{2(1 - \lambda)^2/p + 12\lambda^2/p}}{6\lambda/p} \\
&= \frac{\sqrt{2p}}{6\lambda}(\sqrt{(1 - \lambda)^2 + 6\lambda^2} - (1 - \lambda)) \\
&= \frac{\lambda\sqrt{2p}}{\sqrt{(1 - \lambda)^2 + 6\lambda^2} + (1 - \lambda)}
\end{aligned}$$

and for $\lambda \approx 0$ we have $q^* \approx \lambda^2 p/2 > 0$.

It is clear that the main issue is whether the derivative at $q = 0$ is infinite (as is the case for $\beta < 1$) or finite (as is the for $\beta \geq 1$). In the former case we will have $q^* > 0$ when $\lambda > 0$ and in the latter case we will have $q^* = 0$ for small values of λ.

▶ **Theorem 7.** *For Pareto distribution over intrinsic valuation: (1) for any $\beta \geq 1$ there is a constant $\gamma_\beta > 0$ such that for $\lambda \in [0, \gamma_\beta]$ the optimal seller quality is $q^* = 0$, and (2) for $0 < \beta < 1$, for any $\lambda > 0$ the optimal seller quality is $q^* > 0$.*

Proof. Consider the case that $\beta \geq 1$ we have

$$\frac{\partial}{\partial q} R = \lambda(\frac{q}{p^2})^{\beta-1}((p-q)\frac{\beta}{p^2} - \frac{q}{p^2}) - (1-\lambda)(\frac{1}{2p})^{\beta}$$

Clearly,

$$(p-q)\frac{\beta}{p^2} - \frac{q}{p^2} \leq \frac{\beta}{p}$$

Therefore,

$$\frac{\partial}{\partial q} R \leq \lambda(\frac{q}{p^2})^{\beta-1}\frac{\beta}{p} - (1-\lambda)(\frac{1}{2p})^{\beta}$$

Since $q < p$ (if $q = p$ then $R = 0$) we have

$$\frac{\partial}{\partial q} R < (\frac{1}{p})^{\beta}(\lambda\beta - (1-\lambda)(\frac{1}{2})^{\beta})$$

Therefore, for $\lambda < \frac{1}{1+\beta 2^{\beta}} = \gamma_{\beta}$ we have $\frac{\partial}{\partial q} R < 0$ and hence $q^* = 0$.

Consider the case that $0 < \beta < 1$.

$$\frac{\partial}{\partial q} R = \lambda(p-q)\frac{\beta}{p^2}\frac{p^2}{q}(\frac{q}{p^2})^{\beta} - \lambda(\frac{q}{p^2})^{\beta} - (1-\lambda)(\frac{1}{2p})^{\beta}$$

If we have $q^* > p/2 > 0$ we are done. Otherwise we have

$$\frac{\partial}{\partial q} R > \lambda(p/2)\frac{\beta}{q}(\frac{q}{p^2})^{\beta} - (\frac{1}{2p})^{\beta}(\lambda + (1-\lambda)) = \lambda(p/2)\frac{\beta}{q}(\frac{q}{p^2})^{\beta} - (\frac{1}{2p})^{\beta}$$

For $q < (\lambda\beta)^{1/(1-\beta)}p/2$ we have that $\frac{\partial}{\partial q} R > 0$, and hence the optimal quality q^* is at least $(\lambda\beta)^{1/(1-\beta)}p/2 > 0$. ◀

References

1. Kyle Bagwell and Michael H. Riordan. High and declining prices signal product quality. *The American Economic Review*, 81(1):224–239, 1991.
2. Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.
3. Bernd Fricka and Robert Simmons. The impact of managerial quality on organizational performance: Evidence from german soccer. *Managarial and Decision Economics*, 29:593–600, 2008.
4. Paul J. Gertler and Donald M. Waldman. Quality-adjusted cost functions and policy evaluation in the nursing home industry. *Journal of Political Economy*, 100(6):1232–1256, 1992.
5. Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.
6. Amna Kirmani and Akshay R. Rao. No pain, no gain: A critical review of the literature on signaling unobservable product quality. *Journal of Marketing*, 64(2):66–79, 2000.
7. Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
8. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

9 Ariely Dan Shiv Baba, Carmon Ziv. Placebo effects of marketing actions: Consumers may get what they pay for. *Journal of Marketing Research*, 42:383–393, 2005.

10 Steven M. Shugan. Price-quality relationships. *Advances in Consumer Research*, 11:627–632, 1984.

11 Surendra Rajiv Tridas Mukhopadhyay and Kannan Srinivasan. Information technology impact on process output and quality. *Management Science*, 43(12):1645–1659, 1997.

12 Valarie A. Zeithaml. Consumer perceptions of price, quality, and value: A means-end model and synthesis of evidence. *Journal of Marketing*, 52(3):2–22, 1988.