£ecture$ on Pri¢ing

R. Preston McAfee

California Institute of Technology
Pasadena, CA 91125

First Draft

Table of Contents

1: Introductory Blather

The law of one price is false.

The law of one price – identical objects in a competitive market sell at the same price, or at the same price adjusted for transport costs – has a wonderfully simple proof: for otherwise consumers will shop for the lowest price and the high priced outlets will have no sales. This is persuasive, in an "armchair theorizing" sort of way. But visit a U.S. grocery store and see 50% sales on paper towels, or Coca-Cola, or milk, a sale that will last two or three weeks and then the regular price will return. Nearby stores have sales on different items, tuna or baby food. Nothing changed in these markets – there was no Coca-Cola glut or shift in demand for baby food. The price variation is large, the products standardized, and the stores nearby. Were the law of one price to have empirical force, it should have it in this situation, and it does not.

George Stigler (1961) found significant price variation on what seem like quite standardized items. So price variation has been with us for decades.

A simply bizarre phenomenon is the pricing of code-shared airline seats. Code-shares are flights sold by one airline but operated by another carrier. There are regularly substantial differences in prices on these code shares. An important thing to understand is that the seats being sold are the same. McAfee and te Velde (2006) find 30% price differences sustained for months, even on American Airlines and Alaska Airlines, which share frequent flyer miles on these flights as well. Moreover, since the data comes from Orbitz prices, the channel is the same, although the Orbitz pricing may be a consequence of prices posted on the airline websites. There could be difference in the values customers put on contracting with a particular airline directly even beyond frequent flyer miles – the ability to shift to alternatives might matter, for example, and by dreaming up progressively more elaborate motivations, we might indeed account for the substantial price differences. The important point, however, is that those price differences, on very similar items, are a large and salient feature of the data, and not a minor oddity.

One implication of empirical price variation is that attempts to understand the determinants of price, as if price were the result of the intersection of supply and demand, are doomed to fail unless the attempts account for price variation. Firms charge different prices for the same item to different customers or at different times, and this variation is in fact predictable. These notes are an attempt to organize relevant theories.

1.1: Monopoly Pricing

How does a monopoly choose its price and quantity? Even a monopoly faces a demand curve; price increases will be met with quantity decreases.

As a practical matter, monopolies usually choose prices and demand by customers dictates the quantity sold. However, it is slightly convenient to formulate the problem in the reverse way, with the monopoly choosing the quantity and the price determined by demand; the basis for this convenience is the fact that the cost depends on quantity produced, not the price. Let $p(q)$ be the demand price associated with quantity q, and $c(q)$ be the cost of producing q. The monopoly's profits are

$$\pi = p(q)q - c(q).$$

The monopoly earns the revenue $p(q)q$ and pays the cost $c(q)$. This leads to the first order condition, for the profit-maximizing quantity q_m:

$$0 = \frac{\partial \pi}{\partial q} = p(q_m) + q_m p'(q_m) - c'(q_m).$$

The term $p(q) + qp'(q)$ is the familiar concept called *marginal revenue*. It is the derivative of revenue pq with respect to quantity. It is less than the price ($p(q) + qp'(q) < p(q)$) provided demand slopes down and quantity is positive; at a zero quantity, marginal revenue equals price. A monopolist maximizes profit by choosing the quantity q_m which sets marginal revenue equals marginal cost, and charges the price $p(q_m)$. This maximizes profit because selling an extra unit produces incremental profit equal to the marginal revenue, and costs the marginal cost; if marginal revenue exceeds marginal cost, it is profitable to sell an extra unit, while if marginal revenue is less than marginal cost, it is profitable to reduce sales.

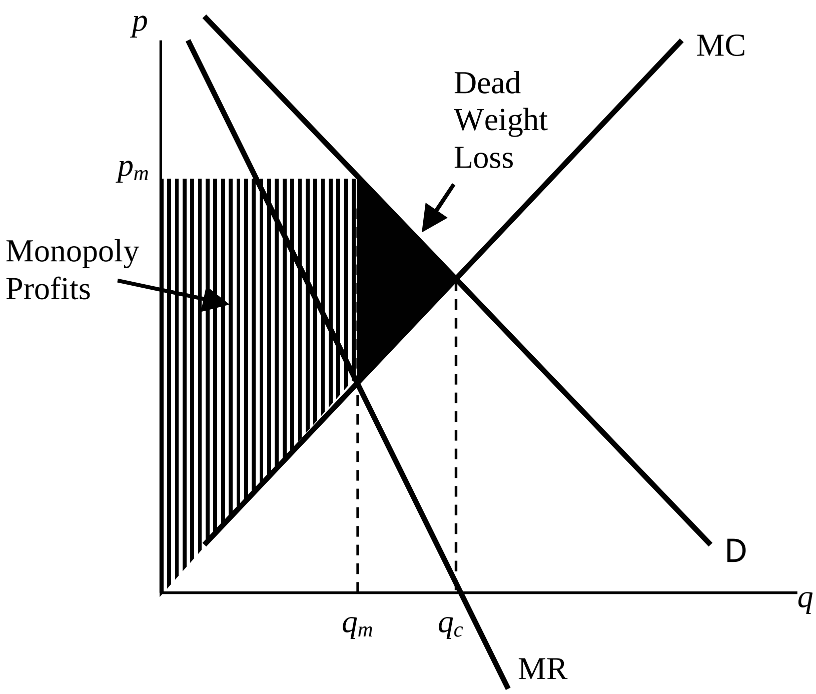

Figure 1: Basic Monopoly Diagram

The basic monopoly diagram is provided in Figure 1. The important features of this figure are:

- marginal revenue lies below the demand curve
- the monopoly quantity equates marginal revenue and marginal cost
- the monopoly price is higher than the marginal cost
- there is a dead weight loss: the higher price of the monopoly prevents some units from being traded that are valued more highly than they cost

- the monopoly profits from the increase in price, and the monopoly profit is shaded
- the intersection of marginal cost and demand corresponds to the competitive outcome (since under competitive conditions supply equals marginal cost)
- the monopoly restricts output and charges a higher price than would prevail under competition.

The monopoly pricing formula can be re-arranged to produce an additional insight. First, recall the elasticity of demand:

$$\varepsilon = -\frac{dq/q}{dp/p} = -\frac{1/q}{p'(q)/p(q)} = \frac{-p(q)}{qp'(q)}.$$

Elasticity is not assumed constant but its dependence on quantity is suppressed for clarity. Re-arranging the monopoly pricing formula,

$$p(q_m) - c'(q_m) = -q_m p'(q_m)$$

and hence

$$\frac{p(q_m) - c'(q_m)}{p(q_m)} = \frac{-q_m p'(q_m)}{p(q_m)} = \frac{1}{\varepsilon}.$$

The left hand side of this equation is known as the *price-cost margin* or *Lerner Index.*[1] The right hand side is one over the elasticity of demand. This formula relates the markup over marginal cost embedded in the price to the elasticity of demand. Because perfect competition forces price to marginal cost, this formula shows that the deviation from perfect competition embodied in proportion of the price which is the markup over marginal cost is one over the elasticity. It is sometimes called an example of an "inverse elasticity rule," although other formulae, and in particular Ramsey pricing, go by that name.

Given free disposal, marginal cost will always be non-negative. If marginal cost is less than zero, the least expensive way to produce a given quantity is to produce more and dispose of the unneeded units. Thus, the price-cost margin is less than one, and as a result, *a monopolist produces in the elastic portion of demand.* One implication of this observation is that if demand is everywhere inelastic (e.g. $p(q) = q^{-a}$ for a>1), the optimal monopoly quantity is essentially zero, and in any event would be no more than one molecule of the product every millennium or so.

The price-cost margin formula above can be re-arranged to obtain:

[1] Abba Lerner, 1903-1982.

$$p(q_m) = \frac{\varepsilon}{\varepsilon - 1} c'(q_m).$$

Thus, a monopolist marks up marginal cost by the factor $\varepsilon / (\varepsilon - 1)$, and this a number exceeding one, since a monopolist only produces at an elasticity greater than one. The formula has been used to justify a "fixed markup policy," which means a company adds a constant percentage markup to its products. This is an ill-advised policy not justified by the formula, because the formula suggests a markup which depends on the demand for the product in question and thus not a fixed markup for all products a company produces.

A monopolist would like to charge customers with a high elasticity a lower price than customers with a low elasticity. But how can the monopoly do so? That is the subject of the next two chapters.

2: Price Discrimination

Price discrimination entails charging distinct customers distinct prices for a good or service. Thus, a "matinee discount" or "early bird special," in which a customer is offered a lower price at a less desirable time, is *not* price discrimination. A student discount or senior citizen's discount, in which the customer gets a discount not because of the nature of the product but based on the identity of the customer, is price discrimination. A common euphemism for price discrimination is "value-based pricing," which means the price is based on the value the customer puts on the good, as opposed to being entirely based on the cost.

Price discrimination comes in two major flavors: direct price discrimination, in which customers are charged based on their identity or based on some observable characteristic of the customer, and indirect price discrimination, in which multiple offers are made to all customers who can then choose the one they like best; if these offers result in distinct prices per unit, price discrimination has occurred. A quantity discount that isn't based on costs is perhaps the most common form of indirect price discrimination.[2] A quantity discount is price discrimination because the customer who chooses to buy many units pays less per unit than the customer who buys one unit. Newspaper coupons, which offer a discount to those who bring the coupon to the store, also represent price discrimination, because they are available to all but only used by some.

There is a certain amount of ambiguity in the definition of price discrimination, most easily seen with transported goods. Steel in Pittsburgh is not literally the same product as steel in Detroit, nor would one expect the two goods to sell at the same price, especially if the steel were produced in Pittsburgh and transported to Detroit. Nevertheless, we may think about the seller engaging in price discrimination if the price "at the mill" differs, that is, if the price net of the transportation cost differs between the two. In a competitive market, the price at the mill would have to be the same, no matter what the final destination. Imagine fifty nearby mills; if the price at the mill differs, all the mills would like to sell to the customers paying the highest price at the mill, which would force equalization. Different prices at the mill, then, is price discrimination. In this case, we have extended the definition of price discrimination to cover the case where the goods were the same at one time, even though transportation has rendered them distinct goods.

Price discrimination provides a rich theory for understanding why prices in the marketplace might show variability. Because distinct customers pay different prices, there will not be a single price at which the market transacts. Our first two chapters are devoted to this entertaining topic.

2.1: Single Purchase

Each consumer demands a single unit, consumers are ranked on a continuum by their type t. Let the distribution of types be F, and index types by their probability $q=F(t)$. Examples of types include age, weight. The willingness to pay of a type t consumer is $p(q)$, which is assumed

[2] Direct and indirect price discrimination replace the cumbersome "degree" notation, in which first degree was perfect price discrimination, second degree was indirect price discrimination, and third degree represents imperfect first degree price discrimination. In this non-mnemonic notation, second isn't even between first and third.

monotone. Given that p is monotone, p is decreasing without loss of generality. Moreover, marginal costs can be subtracted from p; this makes setting marginal cost to zero without loss of generality. The assumption that p is monotone is not without loss of generality because a discriminating monopolist is going to price based on type, which is assumed to be observable; the monotonicity of p insures that conditioning on observable type is equivalent to conditioning on unobservable value.

A non-discriminating monopolist earns $q\,p(q)$; let q_0 maximize profits. A two price discriminating monopolist earns $q_1\,p(q_1) + (q_2 - q_1)p(q_2)$ and abuse notation to let q_1 and q_2 stand for the maximizing arguments. Then

Theorem (Varian 1985): Quantity and welfare (sum of profits and consumer surplus) are higher under price discrimination.

> Proof: Note that welfare depends only on quantity due to the single good purchase by consumers. Thus, it is sufficient to prove that quantity is not lower under price discrimination. Suppose not, that is, suppose $q_2 < q_0$. Then $-p(q_0)\,q_1 > -p(q_2)\,q_1$. Profit maximization for the non-discriminating monopolist insures $p(q_0)\,q_0 \geq p(q_2)\,q_2$. Add these two inequalities to obtain $p(q_0)(q_0 - q_1) > p(q_2)\,(q_2 - q_1)$, which implies
>
> $p(q_1)\,q_1 + p(q_0)(q_0 - q_1) > p(q_1)\,q_1 + p(q_2)\,(q_2 - q_1)$,
>
> which contradicts profit maximization of the two price monopolist. Thus, $q_2 < q_0$ leads to a contradiction. Moreover, welfare is strictly higher if either of the optimizations are strict. In particular, if the non discriminating monopolist has a unique optimal quantity, welfare is higher under discrimination. Q.E.D.

2.2: Selling *n* Goods

This result suggests that price discrimination is invariably a good thing, but that is not a general result. Suppose there are n markets, and demand is given by $x_i(\mathbf{p})$ in market i where $\mathbf{p}=(p_1,\ldots,p_n)$.

$$\pi = \sum_{i=1}^{n} (p_i - mc)x_i(\mathbf{p}).$$

Marginal cost mc is assumed constant. A non-discriminating monopolist charges a constant price p_0 in all n markets. (Demand might be interdependent because the n markets represent distinct goods sold by the same seller, or because of arbitrage across markets. For example, veterinary and human use of medicines has some limited arbitrage. Methyl-methacrylate is used in both dental and industrial uses and historically experienced costly arbitrage. Identical goods sold in many countries are also subject to limited arbitrage. The discriminating monopolist will charge distinct prices p_i in the markets, $i=1,\ldots,n$.

Define the cross-price elasticity of substitution

$$\varepsilon_{ij} = \frac{p_j}{x_i}\frac{dx_i}{dp_j}.$$

Let E be the matrix of elasticities. Note that, if preferences can be expressed as the maximization of a representative consumer, then the consumer maximizes $u(\mathbf{x})$-$\mathbf{px}$, which gives FOC $u'(\mathbf{x}) = \mathbf{p}$, and thus $u''(\mathbf{x})\mathbf{dx} = \mathbf{dp}$. This shows that demand $\mathbf{x}$ has a symmetric derivative, a fact used in the next development.

The first order condition for profit maximization entails

$$0 = \frac{\partial \pi}{\partial p_i} = x_i + \sum_{j=1}^{n}(p_j - mc)\frac{\partial x_j}{\partial p_i} = x_i + \sum_{j=1}^{n}(p_j - mc)\frac{\partial x_i}{\partial p_j} = x_i\left(1 + \sum_{j=1}^{n}\frac{(p_j - mc)}{p_j}\varepsilon_{ij}\right)$$

Let $L_i = \frac{p_i - mc}{p_i}$, and express the first order condition in a matrix format:

$\mathbf{0} = \mathbf{1} + \boldsymbol{E}\,\mathbf{L}$, and thus $\mathbf{L} = -\,\boldsymbol{E}^{-1}\,\mathbf{1}$. This generalizes the well-known one-good case of

$$\frac{p - mc}{p} = -\frac{1}{\varepsilon},$$

where ε is the elasticity of demand (with a minus sign). In the one dimensional case, the price/cost margin (aka the Lerner index) is the inverse of the elasticity. (Usually in the one market case, the minus sign is incorporated into the elasticity definition.) In the *n* market case, the price/cost margins depend on the matrix of elasticities, but still have the simple inverse elasticity form.

In the most frequently encountered version of monopoly pricing, demands are independent, in which case $\boldsymbol{E}$ is a diagonal matrix. The markets are then independent, and

$$\frac{p_i - mc_i}{p_i} = -\frac{1}{\varepsilon_{ii}}.$$

The Robinson-Patman Act of 1933 amended the Clayton act to make price discrimination illegal when the product is sold to intermediaries, rather than final consumers. The prime target of the act was A&P (the Great Atlantic and Pacific Tea Company, which arguably invented the now popular superstore or "category killer"). Does price discrimination increase, or decrease, welfare?

Theorem (Varian, 1985): The change in welfare, ΔW, when a monopolist goes from non-discrimination to discrimination is given by

$$\sum_{i=1}^{n}(p_i - mc)\Delta x_i \le \Delta W \le (p_0 - mc)\sum_{i=1}^{n}\Delta x_i.$$

Proof: Let $\mathbf{p_0} = p_0\mathbf{1}$, be the one-price monopoly price vector, and $\mathbf{p}$ represent the prices of the discriminating monopolist. Let v be the indirect utility function (consumer utility as a function of prices). The indirect utility function is convex, and its derivative is demand (Roy's identity). Therefore,

$$\mathbf{x}(\mathbf{p_0})(\mathbf{p_0} - \mathbf{p}) \le v(\mathbf{p}) - v(\mathbf{p_0}) \le \mathbf{x}(\mathbf{p})(\mathbf{p_0} - \mathbf{p})$$

The change in profits is

$$\Delta\pi = \mathbf{x}(\mathbf{p})(\mathbf{p} - mc\mathbf{1}) - \mathbf{x}(\mathbf{p_0})\mathbf{1}(p_0 - mc)$$

Since the change in welfare is the change in consumer utility plus the change in profits, we have

$$\mathbf{x}(\mathbf{p_0})(\mathbf{p_0} - \mathbf{p}) + \Delta\pi \le \Delta \mathrm{W} \le \mathbf{x}(\mathbf{p})(\mathbf{p_0} - \mathbf{p}) + \Delta\pi,$$

which combines with $\mathbf{\Delta x}=\mathbf{x}(\mathbf{p})-\mathbf{x}(\mathbf{p_0})$ to establish the theorem. Q.E.D.

This theorem has a powerful corollary. If price discrimination causes output to fall, then price discrimination decreases welfare relative to the absence of price discrimination. This result, established by Schmalensee (1981) in a more restricted environment, has a simple proof for the case of independent demands. Consider the case of two markets. Price discrimination's effect on welfare is composed of two terms – a change in total output, and a reallocation of the output across markets. The re-allocation always has a negative impact on welfare, because for any given quantity, welfare is maximized by using a single price, because this single price equalizes the marginal value of the good across markets. Thus, price discrimination has a negative reallocation effect; this can only be overcome if the quantity effect is positive, that is, price discrimination induces a higher output.

Even in the simplest two-market case of linear demand, price discrimination may increase or decrease welfare. To see this, first consider the case where the one-price monopolist serves both markets (Figure 2). In this case, it is straightforward to show that the switch to price discrimination leaves the total output unchanged, so that the only effect is the reallocation, which lowers welfare.

It is straightforward to construct cases where welfare rises under price discrimination. Even in the two-market, linear demand case, if price discrimination opens a new market that is otherwise not served, welfare will rise. Indeed, in this case, price discrimination is a pareto improvement, because the monopolist will leave price in the market served under no price discrimination unchanged. That is, price discrimination lowers price in the unserved market, while leaving price in the market served under no price discrimination unchanged. This outcome is illustrated in Figure 3.

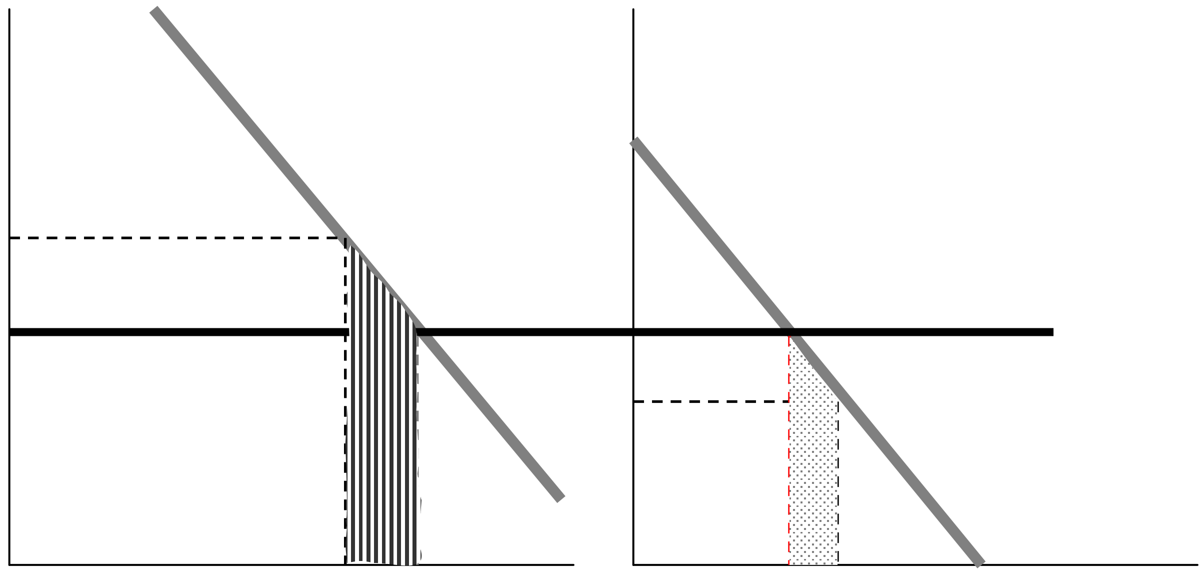

Market 1: Vertical striped area lost by price discrimination

Market 2: Dotted area added by discrimination

Figure 2: Welfare loss from re-allocation under price discrimination.

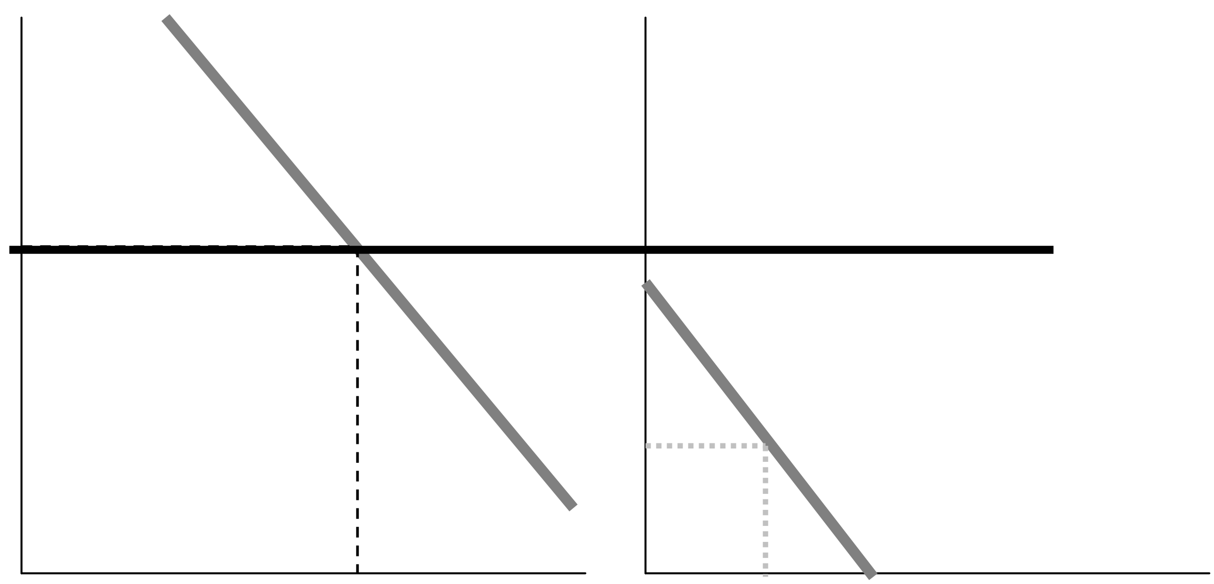

Market 1: Red line indicates no price discrimination outcome.

Market 2: With price discrimination, market 2 is served.

Figure 3: Welfare may rise when price discrimination opens new markets.

2.3: Ramsey Pricing

How should a multi-product or multi-market monopolist be regulated? Ramsey investigated this question. Ramsey pricing is the solution to the problem of maximizing social welfare, subject to a break-even constraint for a monopolist.

In particular, consider the problem

max $u(\mathbf{x})$- $c(\mathbf{x}\bullet\mathbf{1})$ s.t. $\mathbf{p}\bullet\mathbf{x} - \mathrm{c}(\mathbf{x}\bullet\mathbf{1}) \geq \pi_0$.

The symbol $\bullet$ is the standard Euclidean dot product. This formulation permits average costs to be decreasing. Write the Lagrangian

$$\Lambda = u(\mathbf{x}) - c(\mathbf{x}\bullet\mathbf{1}) + \lambda(\mathbf{p}\bullet\mathbf{x} - c(\mathbf{x}\bullet\mathbf{1})) = u(\mathbf{x}) - \mathbf{p}\bullet\mathbf{x} + (1+\lambda)(\mathbf{p}\bullet\mathbf{x} - c(\mathbf{x}\bullet\mathbf{1}))$$

The lagrangian term λ has the interpretation that it is the marginal increase in welfare associated with a decrease in firm profit. Using Roy's identity,

$$0 = \frac{\partial\Lambda}{\partial p_i} = \lambda x_i + (1+\lambda)\sum_{j=1}^{n}(p_j - mc)\frac{\partial x_j}{\partial p_i} = \lambda x_i + (1+\lambda)\sum_{j=1}^{n}(p_j - mc)\frac{\partial x_i}{\partial p_j}$$

$$= \lambda x_i + (1+\lambda)x_i\sum_{j=1}^{n}\frac{p_j - mc}{p_j}\varepsilon_{ij}.$$

Write the first order conditions in vector form, to obtain

$$-\frac{\lambda}{\lambda+1}\mathbf{1} = \mathbf{EL}$$

This equation solves for the general Ramsey price solution:

$$\mathbf{L} = -\frac{\lambda}{\lambda+1}\mathbf{E}^{-1}\mathbf{1}.$$

Note the similarity to monopoly pricing – the welfare optimization problem has the same structural form as the monopoly problem, and moreover the monopoly outcome arises when $\lambda\rightarrow\infty$. Setting $\lambda = 0$ maximizes total welfare and sets price equal to marginal cost in all industries. Such a pricing scheme will give the firm negative profits when average costs are decreasing, because marginal cost is less than average costs.

In general, the Ramsey solution is a mixture of marginal cost pricing and monopoly pricing. That is, the Ramsey solution goes part of the way, but not all of the way, toward monopoly pricing. Such a description is at best an approximation, because elasticities are not constant along the path that connects marginal cost pricing with monopoly pricing, but nevertheless, Ramsey pricing generalizes both with a single formula.

2.4: Arbitrage

Cross-price elasticities can be interpreted as a consequence of arbitrage by individuals. This arbitrage responds in a continuous way to price changes, and thus obtains a “cost per unit” from of arbitrage. For example, suppose leakage from the low priced market to the high priced market costs $\gamma(m)$, where m is the size of the transfer. We let a positive m indicate a transfer from market 1 to market 2, and a negative m the reverse. The marginal cost of transfer is then $\gamma'(m)$. The function γ is assumed convex, with $\gamma'(0) = 0$, which insures that goods flow from the low priced market to the high priced market. Finally, denote the consumer demands in markets 1 and 2 are $q_1(p_1)$ and $q_2(p_2)$. Values in the two markets are assumed independent except for arbitrage effects. The demands facing the seller, x_i, in each market will satisfy:

$$p_2 - p_1 = \gamma'(m),$$

$$x_1 = q_1(p_1) + m, \text{ and}$$

$$x_2 = q_2(p_2) - m.$$[3]

An interesting aspect of these equations is that realized demand is reconcilable with preferences of a single consumer, because

$$\frac{\partial x_i}{\partial p_j} = \frac{\partial x_j}{\partial p_i}.$$

This equation insures that the analysis of the previous section continues to hold. Thus, in particular, arbitrage does not overturn the welfare results already provided, nor does it influence the inverse elasticity results, although the elasticity is a complicated object. Arbitrage does, however, invalidate the independent market model.

It is useful to define the elasticity of arbitrage with respect to price differentials as follows. Since the value of m satisfies $\Delta p = \gamma'(m)$, the elasticity of m with respect to Δp is

$$\eta = \frac{\Delta p}{m}\frac{dm}{d\Delta p} = \frac{\gamma'(m)}{m}\frac{1}{\gamma''(m)} = \frac{\gamma'(m)}{m\gamma''(m)}.$$

---

[3] There is a possibility that the solution involves no purchases in one of the market, if the price in that market is sufficiently high. This possibility is ignored since it will never be part of the profit-maximizing solution; it is always more profitable to set the price slightly below the point that extinguishes sales in one market.

A large value of η means that arbitrage is easy, while a small value means that arbitrage is difficult. In the extreme as $\eta \to 0$, arbitrage becomes impossible. Because $\gamma'(0)=0$ and γ is convex, $\gamma'(m)$ has the same sign as m and η is positive.

It is also useful to recall the elasticities of demand that would prevail in the absence of arbitrage:

$$\varepsilon_i = -\frac{p_i q_i'(p_i)}{q_i(p_i)}.$$

The profits of the seller can be expressed to eliminate the price p_2, which in the process eliminates an implicit dependence on m and makes this dependence explicit.

$$\pi = (p_1 - c)x_1 + (p_2 - c)x_2$$

$$= (p_1 - c)(q_1(p_1) + m) + (p_1 + \gamma'(m) - c)(q_2(p_1 + \gamma'(m)) - m)$$

We can thus view profits as a function of p_1 and m, with p_2 determined by the equation $p_2 = p_1 + \gamma'(m)$. The first order conditions for profit maximization are

$$0 = \frac{\partial \pi}{\partial p_1} = q_1(p_1) + m + (p_1 - c)q_1'(p_1) + q_2(p_1) - m + (p_2 - c)q_2'(p_2)$$

$$= q_1(p_1)\left(1 - \frac{p_1 - c}{p_1}\varepsilon_1\right) + q_2(p_2)\left(1 - \frac{p_2 - c}{p_2}\varepsilon_2\right)$$

and

$$0 = \frac{\partial \pi}{\partial m} = p_1 - c - (p_1 + \gamma'(m) - c) + \gamma''(m)(q_2(p_2) - m + (p_2 - c)q_2'(p_2))$$

$$= -\gamma'(m) - m\gamma''(m) + \gamma''(m)q_2(p_2)\left(1 - \frac{p_2 - c}{p_2}\varepsilon_2\right).$$

Thus,

$$\frac{p_2 - c}{p_2}\varepsilon_2 = 1 - \frac{1}{q_2(p_2)}\frac{\gamma'(m) + m\gamma''(m)}{\gamma''(m)} = 1 - \frac{m}{q_2}(1+\eta).^{4}$$

Similarly,

---

[4] Note that this equation insures that all of good 2's sales cannot come from market 1, for then $m = q_2$ and the price-cost margin would be negative, or $p_2 < c < p_1$, which contradicts $m = q_2 > 0$.

$$q_1(p_1)\left(1-\frac{p_1-c}{p_1}\varepsilon_1\right)=-q_2(p_2)\left(1-\frac{p_2-c}{p_2}\varepsilon_2\right)=-\frac{\gamma'(m)+m\gamma''(m)}{\gamma''(m)}=-m(1+\eta)\,,$$

so

$$\frac{p_1-c}{p_1}\varepsilon_1=1+\frac{m}{q_1}(1+\eta)\,.$$

These rules come in the following form:

$$\frac{p_i-c}{p_i}=\frac{1}{\varepsilon_i}\left(1+\frac{\tau_i}{q_i}(1+\eta)\right),$$

where τ_i is the flow *out* of market *i*. Thus, the monopoly pricing markup formula is adjusted by the following: the fraction of the local demand that flows to the other market, which is negative if there is an inflow, times 1 plus the elasticity of outflow with respect to the price difference. In particular, the markup is increased when the flow is out of the market (low price), and reduced when the flow is into the market. Moreover, the size of this change is proportional to the inverse of the market size. If the outflow is observable and its elasticity can be estimated, these formulas are readily implemented in a real pricing problem plagued by arbitrage.

A heuristic for pricing in the presence of gray markets or leakage across markets is to set "optimal" prices for single markets, and then observe the level of leakage *m*. One observation isn't sufficient to estimate the parameter η, but two points would permit an approximation to the responsiveness to price changes. Once these formula are fit, the adjusted markups are readily computed.

2.5: Preventing Arbitrage

Arbitrage generally limits the ability of sellers to price discriminate, which gives sellers an incentive to limit the ability of customers to engage in arbitrage. How can sellers limit arbitrage?

1. Services

If someone comes to your home, they can charge you a price based on your home's location; you can't resell the service. Airline tickets are non-transferable, which prevents a spot market and severely limits the opportunities of a stand-by passenger.

2. Warranties

A manufacturer may void the warranty if the good is resold; this reduces the value of resale. However, there are limits on the ability of U.S. manufacturers to void warranties in the event of resale and many sellers choose to make warranties transferable.

3. Differentiating products

Brazil switched to ethanol for automotive fuel in the 1970s. To prevent people from drinking this, they added a little gasoline to the fuel. We will explore the ability of manufacturers to differentiate products in Section 3.4:.

4. Transportation costs

For heavy products, like timber, concrete, gravel, and coal, where a major portion of the costs are transportation costs, arbitrage is very expensive. Gasoline refiners charge individual gas stations distinct prices, a form of price discrimination known as zone pricing. Arbitrage is difficult or impossible because it would entail removing the gasoline from an underground storage tank at one station and transporting it to another. Even if this were physically possible, environmental concerns make it impractical.

5. Contracts

I get a free textbook provided I don't sell it. I have sold them anyway in the past; these contracts may be expensive to enforce. There are companies that specialize in shaving the printed language "professor's desk copy" off the resold text. Contracts are often used to limit arbitrage by preventing the buyer obtaining the low price from re-selling.

6. Matching problem

No market for someone who needs what you can buy cheap. Airline tickets were an example, in the 1980s, prior to the internet. For example, passengers going one way and not returning might buy a round trip fare and then attempt to resell the return trip, but it was difficult to find the person needing the reverse one-way. With the internet, finding such people is easier but now tickets are non-transferable.

7. Government

U.S. forces the sale of agricultural products at different prices, with some nations ("most favored nation") enjoying reduced tariffs.

8. Quality

A company may offer different qualities to segregate the market. The high quality sells to high value users, low quality sells to those who can't afford high quality. Such indirect price discrimination is considered in Sections 3.3: and 3.4:.

3: Indirect Price Discrimination

Indirect price discrimination entails using customer choice to accomplish price discrimination. How does this work? Generally, a menu of prices and quantities is offered, and customers choose which offering from the menu suits them best. That results in price discrimination because different customers will rationally choose different bundles. Even though all consumers had the same choice set, some choose the higher price (per unit), though lower total cost. Similarly, with coupons, some choose to pay the time costs of finding and organizing coupons in exchange for a discount, some do not. Coupons work by offering discounts to customers with a low value of time, since those are the customers who will find coupons profitable. Provided customers with a low value of time usually have relatively elastic demand, coupons indirectly offer a discount to the desired target set.

3.1: The Two Type Model

In the two type case, assume there is a consumer L, for low type, with value $v_L(q)$ for quantity q, and H, for high type, with value $v_H(q)$. Both value nothing at zero, so $v_L(0)=v_H(0)=0$. The high type is assumed to have higher demand for every positive quantity $q>0$:

(1) $v'_H(q) \geq v'_L(q)$.

The monopolist offers two quantities q_L and q_H at prices R_L and R_H, respectively, targeted to the consumers L and H. In order for consumers to agree to purchase, two conditions known as individual rationality conditions must be satisfied

(IR_L) $v_L(q_L)$-R_L
(IR_H) $v_H(q_H)-R_H \geq 0$.

Note that, rather than offer a plan in which the consumers don't participate, the monopolist could just as well offer $(q,R)=(0,0)$ and get the same outcome, in which case IR is satisfied. In addition, the monopolist must offer plans constructed so that L chooses (q_L,R_L) and type H chooses (q_H,R_H). The conditions governing these plans are called incentive compatibility conditions and are mathematically formulated as follows.

(IC_L) $v_L(q_L)-R_L \geq v_L(q_H)-R_H$

(IC_H) $v_H(q_H)-R_H \geq v_H(q_L)-R_L$.

The condition (IC_L) merely states that the utility the L consumer gets from purchasing the L plan is at least as great as if the L consumer purchases the H plan. Note that, if the monopolist had designed the plan so that the L consumer chose to purchase the H plan, he could have just as well offered the H plan to the L consumer in the first place, so that IC_L would hold. Thus, IC_L can be considered a constraint on the monopolist, and is without loss of generality. The IC_H constraint is analogous.

The monopolist is assumed to have a constant marginal cost c, and to maximize profit $R_L + R_H - c(q_L + q_H)$.

The analysis of the monopolist's behavior is performed by a series of claims, which will simplify the problem until a solution is obvious.

Claim 1: $q_L \leq q_H$.

Proof: Rearrange IC_L and IC_H to obtain

$$v_H(q_H) - v_H(q_L) \geq R_H - R_L \geq v_L(q_H) - v_L(q_L).$$

This gives

$$\int_{q_L}^{q_H} v_H'(q)\, dq = v_H(q_H) - v_H(q_L) \geq v_L(q_H) - v_L(q_L) = \int_{q_L}^{q_H} v_L'(q)\, dq,$$

or,

$$\int_{q_L}^{q_H} v_H'(q) - v_L'(q)\, dq \geq 0,$$

from which (1) proves the claim.

Claim 2: IR_H can be ignored. That is, IC_H and IR_L imply IR_H.

Proof: Using first IC_H then IR_L, note that

$$v_H(q_H) - R_H \geq v_H(q_L) - R_L = \int_0^{q_L} v_H'(q)\, dq - R_L \geq \int_0^{q_L} v_L'(q) dq - R_L = v_L(q_L) - R_L \geq 0.$$

Thus, if IC_H and IR_L are satisfied, then IR_H is automatically satisfied, and can be ignored.

Claim 3: IC_H is satisfied with equality at the monopolist's profit maximization.

Proof: Suppose not. Then the monopolist can increase R_H up to the point where IC_H is satisfied with equality, without violating either IR_L or IC_L. Since this increases revenue, the monopolist would do so, contradicting the assumption that the monopolist had maximized profit.

Claim 4: IC_L is redundant given $q_L \leq q_H$.

Proof: This follows from the fact that IC_H is satisfied with equality, so the quantity restriction (see Claim 1) gives

$$R_H - R_L = v_H(q_H) - v_H(q_L) \geq v_L(q_H) - v_L(q_L)$$

which implies IC_L.

Claim 5: IR_L holds with equality.

> Proof: Otherwise the monopolist could raise both R_L and R_H by the same amount, without violating the constraints.

Claims 3 and 4 let us express the monopolist's objective function in terms of the quantities, merely by using the constraints that hold with equality. That is,

$$R_L + R_H - c(q_L+q_H) = 2v_L(q_L) + v_H(q_H) - v_L(q_H) - c(q_L+q_H).$$

This gives the first order conditions

$$0 = v'_H(q_H) - c,$$

and

$$0 = 2v'_L(q_L) - v'_H(q_L) - c.$$

The second equation may not be satisfiable, and in fact, if the demand of the high type is twice or more the demand of the low type, that is, $v'_H(q) > 2v'_L(q),$ then the monopolist's optimal quantity q_L is 0, and the low type is shut out of the market.

We can deduce the following insights from these equations and the IR and IC constraints.

1. The high type gets the "efficient" quantity (i.e. the quantity that a benevolent social planner would award him.

2. The low type gets strictly less than the efficient quantity.

3. The high type has a positive consumer surplus, that is, $v_H(q_H)$-R_H>0, unless q_L=0.

4. The low type gets zero consumer surplus.

How general are the insights generated from this model? It is a straightforward extension to have unequal numbers of the two types. But two types itself is special, and the next section explores a monopolist facing infinitely many types.

3.2: The Continuum Model

Suppose consumers have utility $v(q,t) - p$, where t is the type in [0,1] with density $f(t)$, q is quantity and p is the payment made. The monopolist will place a aggregate charge $R(q)$ for the purchase of q. What should the schedule of prices $R(q)$ look like?

Define the shadow price $p(q,t) = v_q(q,t)$, which gives the demand curve of the type t.

Assume $p_t(q,t) > 0$, that is, higher types have higher demands, and that $v(0,t) = 0$.

We will look for a function $q^*(t)$ so that a type t agent purchases $q^*(t)$. Any candidate function $q(t)$ must satisfy

(IC) $v(q(s),t)\text{-}R(q(s)) \leq v(q(t),t) - R(q(t)) = \pi(t)$

yielding the first order condition

$$v_q(q(t),t) - R'(q(t)) = 0.$$

and thus, by the envelope theorem,

$$\pi'(t) = v_t(q(t),t).$$

As before, the individual rationality constraint requires

(IR) $\pi(t)\geq 0$.

However, since π is non-decreasing (as $v_t(q,t) = \int_0^q v_{qt}(x,t)dx = \int_0^q p_t(x,t)dx \geq 0$), IR is equivalent to $\pi(0)\geq 0$.

Therefore,

$$\int_0^1 \pi(t) f(t)dt = -\pi(t)(1-F(t)) \Big|_0^1 + \int_0^1 \pi'(t)(1-F(t))dt$$

$$-\pi(0) + \int_0^1 v_t(q(t),t)(1-F(t))dt$$

Consequently, the monopolist's profit can be expressed as:

$$\int_0^1 R(q(t)) - cq(t)) f(t)dt = \int_0^1 (v(q(t),t) - \pi(t) - cq(t)) f(t)dt$$

$$= - \pi(0) + \int_0^1 \left(v(q(t),t) - \frac{1-F(t)}{f(t)} v_t(q(t),t) - cq(t) \right) f(t)dt.$$

Because of IR, the monopolist will set π(0)=0; otherwise he charge all buyers the additional amount π(0), increasing his profit and still satisfying IR and IC. Maximizing point-wise gives:

(2) $$p(q^*(t),t) - \frac{1-F(t)}{f(t)} p_t(q^*(t),t) - c = 0.$$

Lemma (Necessity and Sufficiency): The IC constraint holds if and only if the first order condition for the buyer's maximization holds and q is non-decreasing.

> Proof: Let $u(s,t)=v(q(s),t) - R(q(s))$, which is what a type t agent gets if he buys the quantity slated for type s. Then IC can be written
>
> $u(s,t) \leq u(t,t)$.
>
> Denote partial derivatives with subscripts. Necessarily, $u_1(t,t)=0$ and $u_{11}(t,t)\leq 0$. Totally differentiating the first gives $u_{11}(t,t) + u_{12}(t,t)=0$, so the second order condition can be rewritten $u_{12}(t,t)\geq 0$. Therefore, necessarily,
>
> $0 \leq v_{qt}(q(t),t)\, q'(t)$,
>
> which forces q non-decreasing, since $v_{qt}=p_t>0$. Now turn to sufficiency. Note that, if q is non-decreasing, then u_{12} is everywhere nonnegative. Thus, for $s<t$, $u_1(s,t) \leq u_1(s,s) = 0$, and for $s>t$, $u_1(s,t) \geq u_1(s,s) = 0$. Thus, u is increasing in s for $s<t$, and decreasing in s for $s>t$, and therefore u is maximized at $s=t$, and IC holds. Q.E.D.

Thus, $q^{*\prime}(t) \geq 0$ is both necessary and sufficient for the solution to

> $R'(q^*(\mathrm{t})) = p(q^*(t),t)$
>
> $R(q^*(0)) = v(q^*(0),0)$

to maximize the monopolist's profit, where q^* is given by (2). This defines the optimal R.

Observations:

(1) The highest type consumer gets the efficient quantity, in that price $p(q^*(1),1) = c$, which is marginal cost

(2) Those with greater demand (high t's, since $p_t>0$) obtain at least as much of the good, and sometimes more, than those with lower demand.

(3) All agents except the highest type get less than the efficient quantity

This follows from $p(q^*(t),t) - c = \frac{1-F(t)}{f(t)} p_t(q^*(t),t) > 0$,

since

$$v_t(q,t) = \int_0^q v_{qt}(r,t)dr > 0.$$

(4) If the optimal quantity is decreasing in some neighborhood, then a flat spot results from the optimization and an interval of types are treated equally. This is called pooling.

The monopolist's solution may be implemented using a nonlinear price schedule. Under some circumstances, it may be implemented using a menu of linear price schedules, that is, offering lower marginal costs, at a higher fixed cost, much like phone companies do.

The solution can be interpreted according to the elasticity formula already given. Let $y = 1 - F(t)$ represent the number of consumers willing to buy $q(t)$ at price $p(q(t),t)$. Note that

$$\frac{p(q(t),t)-c}{p(q(t),t)} = \frac{1-F(t)}{f(t)} \frac{p_t(q(t),t)}{p(q(t),t)} = \frac{\frac{\partial}{\partial t}\log(p(q,t))}{-\frac{\partial}{\partial t}\log(1-F(t))} = -\frac{1}{\frac{p}{y}\frac{dy}{dp}}.$$

Thus, we have the usual inverse elasticity rule holding, even in the case of price discrimination with a continuum of types.

3.3: Quality Premia

We have a quality premium model without doing any work, by merely reinterpreting the two tye quantity discounts model of section 3.1:. Suppose the monopolist faces two types of consumers, L and H. The monopolist has at his disposal a range of qualities to offer. Both types value higher quality more, but the H type values an increase in quality more than the L type, that is, $v_L'(q) < v_H'(q)$. In this case, the monopolist will offer two qualities, one high and one low. The high quality good will be efficient, i.e. sets the marginal value of quality to the marginal cost. The low quality, however, will be worse than efficient. That is, the monopolist will intentionally make the low quality good worse, so as to be able to charge more for the high quality good.

The case where c=0 is especially interesting, because this is the case in which quality is free, say, up to an upper bound $\bar{q}$. One can imagine that the monopolist only produces one good, and, at no cost, can make it lower quality, say, by hitting it with a hammer. In this case, the monopolist will still offer two qualities, that is, the monopolist will intentionally damage a portion of the goods he sells, so as to be able to segment the market. It is worth thinking about whether this is what goes on at outlet malls and stores like Sam's and The Price Club. Manufacturers create inconvenient sizes of products, or locate outlets at distant (although not necessary less expensive) locations, in order to be able to charge less to the more price sensitive segment of their market.

It is a straightforward exercise to adapt the two type model so that it is more costly to offer lower quality, that is, the manufacturer has to take an existing product and pay to have it damaged. The only thing that is needed is to replace the cost $c(q_L+q_H)$ with $c_L(q_L) + c_H(q_H)$, where $c_L > c_H$. In this case, the manufacturer may still offer the low quality, that is, pay extra to have some of the goods damaged. The objective of this action is the same, that is, to deter the high demand types from buying the low quality, by reducing the low quality below efficient levels. A more thorough discussion of damaged goods is contained in the next section.

3.4: Damaged Goods

Manufacturers intentionally damage some of the products they sell to make them less useful, so that they can be sold as a discount. For example, a Saturday night stay-over restriction, which is used by airlines to justify discounts on seats, does not prevent travel at peak times, but instead merely hampers the ability of a traveler to select an outbound and return flight at the times they might desire. This injures the product for some customers more than for others, which permits the airline to charge more for the fare with no such restriction, even when the same seats are occupied.

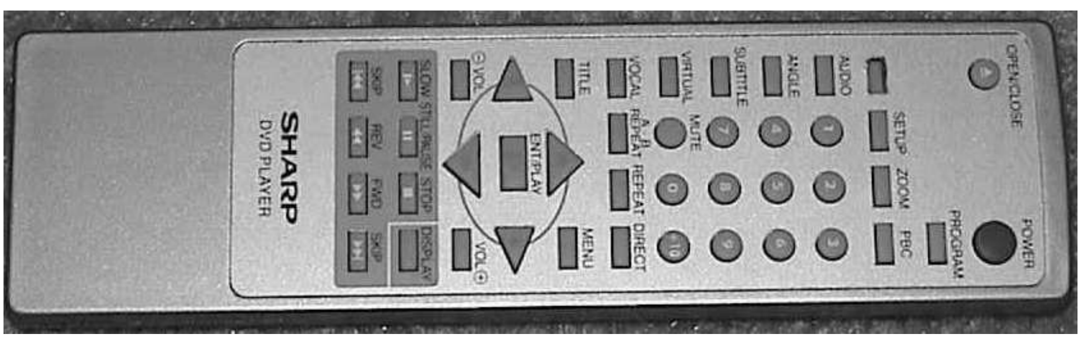

Figure 4: Hacked Remote Control of the DV740U (Courtesy of Area 450). Note extra button in upper right hand corner.

Figure 4 illustrates a damaged product. In this case, the product is a DVD player, where a useful function was suppressed in one version of the product by the artifice of hiding the button that would evoke it, and an industrious user has cut a hole in the top of the remote to access the button. (See McAfee (2006) for details.)

To develop a theory of damaged products, consider the situation where customers who value a good at v value the damaged good at $\lambda(v)$, where $\lambda(0)=0$, λ is increasing, and $\lambda' < 1$. These assumptions seem reasonable for many situations, although don't cover cases where high value customers would rather have nothing than the damaged product.

Suppose the price of the high quality product is p_H and the price of the low quality product is $p_L < p_H$. Consumers with value $v<x_L$ will buy nothing, consumers with value $x_L<v<x_H$ will buy the low quality product, and consumers with $v>x_H$ will buy the high quality product. These values satisfy

$$p_L = \lambda(x_L),$$

and

$$x_H - \lambda(x_H) = p_H - p_L .$$

Let F be the distribution of valuations, and suppose the marginal cost of production is constant at c. We can express profits as

$$\begin{aligned}\pi &= (1-F(x_H))(p_H - c) + (F(x_H) - F(x_L))(p_L - c) \\ &= (1-F(x_H))(p_H - c) + [(1-F(x_L)) - (1-F(x_H))](p_L - c) \\ &= (1-F(x_H))(p_H - p_L) + (1-F(x_L))(p_L - c) \\ &= (1-F(x_H))(x_H - \lambda(x_H)) + (1-F(x_L))(\lambda(x_L) - c) .\end{aligned}$$

It is useful to take a digression to consider what marginal revenue is for this structure. We have two separate marginal revenues: the marginal revenue for the regular good and the marginal revenue for the low quality good, and we will assume both are decreasing in quantity. The marginal revenue for the main good is the derivative of total revenue with respect to quantity, which in our notation is

$$MR_H = \frac{\frac{d}{dp}pq}{\frac{dq}{dp}} = \frac{\frac{d}{dp}p(1-F(p))}{\frac{d}{dp}(1-F(p))} = \frac{-pf(p) + (1-F(p))}{-f(p)} = p - \frac{(1-F(p))}{f(p)}$$

Recalling that marginal revenue is decreasing in quantity if and only if it is increasing in price, decreasing marginal revenue is tantamount to assuming $p - \frac{(1-F(p))}{f(p)}$ is increasing in price.

Similarly, the quantity demanded of the crimped good, were the regular good not for sale, is given by $1 - F(\lambda^{-1}(p))$ for price p; the marginal revenue for the crimped product is

$$MR_L = \frac{\frac{d}{dp}pq}{\frac{dq}{dp}} = \frac{\frac{d}{dp}p(1-F(\lambda^{-1}(p)))}{\frac{d}{dp}(1-F(\lambda^{-1}(p)))} = \frac{-pf(\lambda^{-1}(p))\lambda^{-1\prime}(p) + (1-F(p))}{-f(\lambda^{-1}(p))\lambda^{-1\prime}(p)}$$

$$= p - \frac{(1 - F(\lambda^{-1}(p)))}{f(\lambda^{-1}(p))} \lambda'(\lambda^{-1}(p)) .$$

Since λ is increasing, marginal revenue for the low quality good is decreasing if and only if $\lambda(p) - \frac{1 - F(p)}{f(p)} \lambda'(p)$ is increasing.

The main theorem, proved in McAfee (2006), demonstrates that if $\frac{\lambda(v) - c}{v - c}$ is increasing for all $v \geq \lambda^{-1}(c)$, then it is unprofitable to offer the crimped good. In contrast, if $\frac{\lambda(v) - c}{v - c}$ is decreasing around the monopoly price for the regular good, is decreasing, then it does pay to offer the crimped good. The results are independent of the distribution *F* beyond the requirement that the marginal revenues are decreasing. To see why this is true, note that $\frac{\lambda(v) - c}{v - c}$ is increasing when $(v - c)\lambda'(v) - (\lambda(v) - c) > 0$. Then

$MR_L(\lambda(p)) - c = 0$

if and only if $\lambda(p) - c - \frac{1 - F(p)}{f(p)} \lambda'(p) = 0$

if and only if $\frac{\lambda(p) - c}{\lambda'(p)} = \frac{1 - F(p)}{f(p)}$

if and only if $MR_H - c = p - \frac{(1 - F(p))}{f(p)} - c = p - c - \frac{\lambda(p) - c}{\lambda'(p)} = \frac{1}{\lambda'(p)}\left(p - c - \lambda(p) - c\right) > 0$.

Thus, at the price at which $MR_L = c$, $MR_H > 0$, which means that it is profitable to have a lower price on *H* than the quality adjusted price on *L*. But this kills the market for low quality goods. Another way to view this result is that, at the monopoly price for *H*, the marginal revenue on *L* is negative, so that there are no sales of *L*. Conversely, when $\frac{\lambda(v) - c}{v - c}$ is decreasing at the monopoly price of *H*, the marginal revenue on *L* is positive, so that it pays to sell a bit more of the low quality good. Thus offering both products is optimal.

The results make intuitive sense. When $\frac{\lambda(v) - c}{v - c}$ is increasing, high value consumers value the crimped product relatively more than low value consumers, so offering a crimped version doesn't work very well as a price discrimination tool. In contrast, when $\frac{\lambda(v) - c}{v - c}$ is decreasing, high value customers don't like the crimped version very much, so it can be sold to low value types without much of a price cut to high value customers.

In addition, it is readily proved that the "ideal" damage to the good offers very high value to low value customers and low value to high value customers – $\lambda(x_L) = x_L = \lambda(x_H)$. Thus, the trick to crimping products is to crimp only the parts that high value customers want and not the part that low value customers want.

The feature removed using the Sharp remote control was the ability to play European DVDs and output the signal to a US television. Was this a sensible feature to remove? Probably, because people who have both European DVDs and US televisions, which are incompatible, are probably world travelers, and have relatively high incomes and willingness to pay. On the other hand, most regular customers would never use the feature.

In contrast, Sony's minidiscs came in two versions, 60 minute and 74 minute. These differed by software instructions that prohibited writing on part of the disc. (McAfee and Deneckere, 1996.) Was this a profitable crimping strategy? Probably not, because most customers would likely value the discs approximately proportional to their length, which makes λ linear, which in turn implies that $\frac{\lambda(v)-c}{v-c}$ is increasing. In this situation, the crimped product is not profitable to offer.

3.5: Tie-ins

"Buy a suit and get an electric drill."
-Detroit TV Ad, 1981

"Shoe: Buy one, get one free".
-South Carolina Billboard, 1987

Tie-ins arise whenever a manufacturer requires the purchase of one product in order to purchase another product. Thus, if an automobile manufacturer required you to use their parts when you had the car serviced, a tie-in would have occurred.

Reasons for Tie-ins

1. Lower Cost

Tie-ins may be lower cost because they save on packaging – e.g. left and right shoes might as well be sold in the same box, which saves not only on cardboard but on organization and mis-matched shoes. For this reason, cars come with tires and radios and the like. In addition, tie-ins may save on sorting costs; the most famous example is probably de Beers, which sells similar grade diamonds in a package that can't be split apart. A similar phenomenon probably accounts for potatoes, oranges and the like sold in packages in grocery stores.

2. Evade price controls

Bundling a price-controlled good with an uncontrolled good can help circumvent price controls, although regulators usually see through such artifice.

3. Offer Secret Price Cuts

Firms in a cartel may want to lower price without their competitors knowing; by bundling a good with another good, it is possible to conceal a price cut.

4. Assure Quality

Kodak tried to sell film with development included, which prevents consumers from blaming Kodak for bad film development by a third party; Kodak was prevented by an antitrust suit from engaging in this practice, which was viewed by the court as an attempt to extend monopoly power in film into film development.

5. Price Discriminate

The general idea is to tie a product with a variable volume to another product, and use the variable volume product to charge heavy users a higher price. This is sometimes known as a Gillette strategy, after their famous practice of selling razors cheaply and making profits on blades. That way, a consumer who uses the razor infrequently or is just testing it doesn't pay much, while a consumer who uses it a lot pays a lot. This was IBM's strategy with business machines and punch cards; IBM rented the machines inexpensively (relative to cost) and then charged a lot – more than ten times cost – for cards. Heavy users, who are probably the high value users, used a lot of cards and thus paid a lot more.

The courts enjoined IBM from this practice, and there are a series of ruling that have a common theme: once a product has been purchased, the manufacturer loses control, and can't force a consumer to do anything with it. For example, a manufacturer can't condition a warranty on the use of the manufacturer's parts if other parts of equal quality are available. The manufacturer can condition the warranty on the use of parts of adequate quality. If your car engine blows up because you put a bad water pump from another manufacturer on it, then this can void the warranty. If you can show, however, that the water pump was of equal or better quality than the original car manufacturer's pump, then your warranty will still be valid.

Many of the lawsuits concerning tie-ins are between franchisees and franchisors, and concern whether a franchisee has to by the franchisor's products (e.g. does a McDonald's franchise have to buy the napkins and coffee stirrers sold by the McDonald's corporation?). This occurs because of the main things a franchise offers is a nationwide quality standard. Thus, people go into McDonald's when travelling in California because they are familiar with the quality in Texas, and expect that quality to be the same. An individual franchise, however, often has an incentive to cut quality (e.g. filthy bathrooms) because it saves on costs, and most of the impact is felt by other outlets (e.g. a McDonald's on an interstate highway gets little repeat business, so it loses few customers because of low quality, but people who do go there are less likely to go into other McDonalds. For this reason, McDonald's polices the quality very carefully). This is only half of the story. The franchisor has an incentive, once the business is not growing rapidly, to try to increase the prices of inputs to the franchisees, who have sunk a large investment in the business and are unlikely to go bankrupt because of an input price increase. Thus, once the business has become mature, a franchisor that requires the use of its own inputs might raise the price to the franchisees. So the lawsuits usually revolve on (i) is the franchisee trying to cut quality by using cheaper inputs than the franchisor's own brand, or (ii) is the franchisor overcharging for the inputs? In most of these cases, experts are brought in to compare the quality of the franchisor's product and the competing brand that the franchise wants to buy.

Generally, a franchisor has the right to insist on a minimum quality, but does not have the right to insist on the use of its own products, if products of equal quality are available.

Bundling, on the other hand, is generally legal. However, offering a lower price for two products together than the sum of the individual prices, called mixed bundling (pure bundling is when you don't offer the two products separately as well as in a bundle, e.g. they don't sell new cars without tires, so cars + tires are pure bundling, but they do sell new cars without radios, new cars with radios, and radios, so cars + radios are mixed bundling), can be illegal if it is found to be price discrimination (see Robinson-Patman Act).

It turns out that, under reasonable specifications of preferences, a monopolist always prefers mixed bundling to no bundling, that is, the monopolist will always set a price for the bundle lower than the sum of the individual prices, as we show in the next subsection.

3.6: Bundling

Former long distance company AT&T reported that one of the most effective marketing tools that it used to sell long distance telephony was a discount available with an oil change at Jiffy Lube (McAfee, 2002). This is a peculiar discount, because oil changes seemingly have little to do with telephony. The remarkable fact is that discounts for the combined purchase of unrelated products are in fact profitable for sellers. This fact was illustrated by a numerical example in Adams and Yellen (1976) and proved in general by McAfee, McMillan and Whinston (1989). To see why bundling is generally profitable, consider two goods, labeled 1 and 2. Any given consumer has a value v_1 for good 1 and v_2 for good 2, and the value of consuming both is just the sum of the individual values. In this formulation, consumers buy at most one unit of each good. This makes the values enter utility independently. In addition, we assume the distribution of consumer valuations for the two goods is independent. Thus, if V_i is the random variable representing the value of good i, we have

$$\text{Prob}\{V_1 \le v_1 \ \& \ V_2 \le v_2\} = F_1(v_1)F_2(v_2)$$

To simplify the analysis, it is useful to assume the cumulative distribution functions F_i have continuous densities f_i with support that is an interval.

Given prices p_1 and p_2 for goods 1 and 2 respectively, and a price p_B for purchasing both, a consumer with values v_1 and v_2 has a choice of purchasing nothing, purchasing 1 only, purchasing 2 only, or purchasing the bundle. This produces the following utilities

The consumer buys good 1 when $v_1 \ge p_1$, $v_1 - p_1 \ge v_2 - p_2$, and $v_1 - p_1 \ge v_1 + v_2 - p_B$. For the case when $p_B \le p_1 + p_2$, the various purchase regions are illustrated in Figure 5. The consumer preferring good 1 to buying the bundle satisfies $v_1 - p_1 \ge v_1 + v_2 - p_B$ or $p_B - p_1 \ge v_2$, which automatically insures $p_2 \ge v_2$ since by hypothesis, $p_2 \ge p_B - p_1$. Thus, the "Buy Good 1 Only" region is determined by the two conditions $v_1 \ge p_1$ and $p_B - p_1 \ge v_2$, both straight lines illustrated in the figure. The other cases require similar justification.

Action	Utility
Buy Nothing	0
Buy Good 1	$v_1 - p_1$
Buy Good 2	$v_2 - p_2$
Buy Both	$v_1 + v_2 - p_B$

Figure 5 shows that a firm with constant marginal cost's profit of offering the prices p_1, p_2 and p_B, with $p_B \le p_1 + p_2$, is

$$\pi = (p_1 - c_1)(1 - F_1(p_1))F_2(p_B - p_1) + (p_2 - c_2)(1 - F_2(p_2))F_1(p_B - p_2)$$

$$+ (p_B - c_1 - c_2)\left(\int_{p_B - p_2}^{p_1} f_1(x)(1 - F_2(p_B - x))dx + (1 - F_1(p_1))(1 - F_2(p_B - p_1)) \right)$$

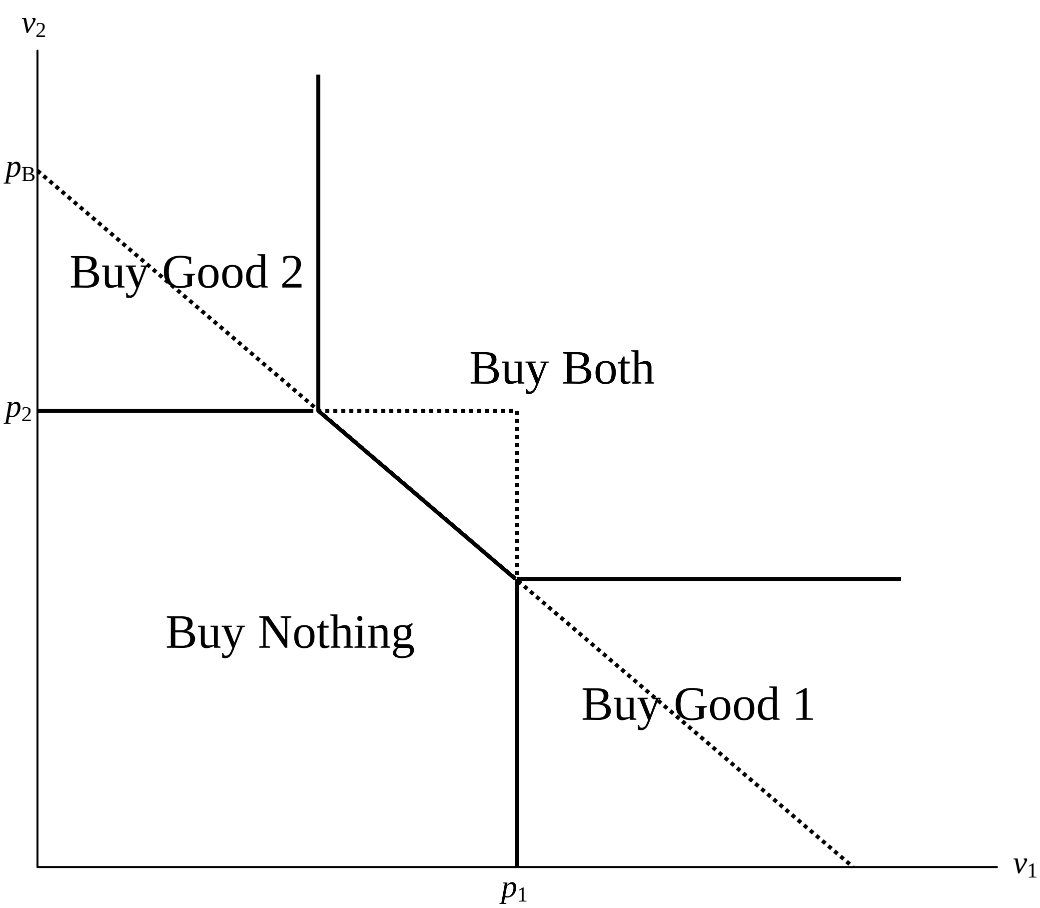

Figure 5: Optimal Consumer Choice

To see if bundling is profitable, we hypothesize $p_B = p_1 + p_2 - \varepsilon$ for a small $\varepsilon>0$, and p_1 and p_2 set to maximize profit. In this case,

$$\pi = (p_1 - c_1)(1 - F_1(p_1))F_2(p_2 - \varepsilon) + (p_2 - c_2)(1 - F_2(p_2))F_1(p_1 - \varepsilon))$$

$$+ (p_1 + p_2 - \varepsilon - c_1 - c_2)\left(\int_{p_1-\varepsilon}^{p_1} f_1(x)(1 - F_2(p_1 + p_2 - \varepsilon - x))dx + (1 - F_1(p_1))(1 - F_2(p_2 - \varepsilon))\right)$$

$$= (p_1 - c_1)\left[(1 - F_1(p_1)) + \left(\int_{p_1-\varepsilon}^{p_1} f_1(x)(1 - F_2(p_1 + p_2 - \varepsilon - x))dx\right)\right]$$

$$+ (p_2 - c_2)\left[(1 - F_2(p_2)) + \left(\int_{p_2-\varepsilon}^{p_2} f_2(x)(1 - F_1(p_1 + p_2 - \varepsilon - x))dx\right)\right]$$

$$- \varepsilon\left(\int_{p_1-\varepsilon}^{p_1} f_1(x)(1 - F_2(p_1 + p_2 - \varepsilon - x))dx + (1 - F_1(p_1))(1 - F_2(p_2 - \varepsilon))\right)$$

Thus,

$$\left.\frac{\partial \pi}{\partial \varepsilon}\right|_{\varepsilon=0} = (p_1 - c_1) f_1(p_1)(1 - F_2(p_2)) + (p_2 - c_2) f_2(p_2)(1 - F_1(p_1))$$

$$- (1 - F_1(p_1))(1 - F_2(p_2))$$

It is straightforward to establish that the profit-maximizing level of p_1 and p_2 entail

$$0 = \left.\frac{\partial \pi}{\partial p_1}\right|_{\varepsilon=0} = (p_1 - c_1) f_1(p_1)(1 - F_2(p_2)) - (1 - F_1(p_1))(1 - F_2(p_2))$$

and

$$0 = \left.\frac{\partial \pi}{\partial p_2}\right|_{\varepsilon=0} = (p_2 - c_2) f_2(p_2)(1 - F_1(p_1)) - (1 - F_1(p_1))(1 - F_2(p_2)).$$

Thus,

$$\left.\frac{\partial\pi}{\partial\varepsilon}\right|_{\varepsilon=0} = (p_1 - c_1) f_1(p_1)(1 - F_2(p_2)) > 0.$$

That is, a small reduction in the price of the bundle, below the sum of the component prices, increases profits.

Why is bundling profitable, even though the goods are unrelated to each other? Bundling makes a price reduction do double-duty. A reduction on the price of the bundle encourages customers who are buying good 1 and are on the margin of buying good 2 to buy both goods, but at the cost of a price reduction for those buying both goods. This is an effect exactly equal to the effect of a reduction in the price of good 2 for the set of people who have a high value of good 1, which is the same effect as a reduction in the price of good 2, due to independence of goods 1 and 2. Thus, to the first order, this effect is zero. But in addition, the reduction in the price of the bundle also encourages the customers on the margin of buying good 1 to buy good 1. Since we have already accounted for the price reduction to those not on the margin, these extra sales are pure profit!

Thus, it is always profitable to engage in bundling of arbitrary unrelated items like Jiffy Lube and AT&T long distance.

4: Peak Load Pricing

Consider a firm that experiences two kinds of costs – a capacity cost and a marginal cost. How should capacity be priced? This issue is applicable to a wide variety of industries, including pipelines, airlines, telephone networks, construction, electricity, highways, and the internet.

We start our investigation with a consideration of a competitive industry. In this case, a very large number of firms each bring a small amount of capacity to the market. A firm incurs a capacity charge (fixed cost) β and a marginal cost *mc*. All firms supply, unless that would drive prices below *mc*, that is, the supply is either the entire capacity or the quantity that makes price equal to marginal cost. This is the usual textbook competitive supply, because average variable cost is by hypothesis equal to marginal cost.

Let the price realized in the market come in the form $p(Q, s)$, where Q is the total quantity supplied and s is the state of nature that will determine the demand. A firm's profits are

$$\pi = E_s \max\{0, p(Q,s) - mc\} - \beta$$

Free entry entails zero profits, and thus the competitive quantity is that quantity that makes $E_s \max\{0, p(Q,s) - mc\} = \beta$. Note that there must be a positive probability that the industry is capacity constrained, and it is during these constrained times that a competitive industry recoups its capacity costs; when the industry has adequate capacity to set price equal to marginal cost, it covers its variable costs but loses money overall.

Is this competitive solution efficient? The answer of course yes; there are no interferences to the market efficiency such as externalities, public goods, taxes, asymmetries of information or monopoly. To see the efficiency of the industry, note that the gains from trade from selling $q(s)$ in state s is

$$W = E_s \int_0^{q(s)} p(x,s)dx - \beta \max_s \{q(s)\} - mc\, q(s) = E_s \int_0^{q(s)} (p(x,s) - mc)dx - \beta \max_s \{q(s)\}$$

Let $Q = \max_s \{q(s)\}$, and note that, as part of the maximization of W, if $q(s) < Q$, then

$$\frac{dW}{dq} = p(q,s) - mc\,.$$

Thus, we obtain that maximizing W entails $p(q(s),s)=mc$ whenever $q(s) < Q$, and thus that either $p(q(s),s)=mc$ or $q(s)=Q$. This allows us to rewrite W as

$$W = E_s \int_0^{q(s)} (p(x,s) - mc)dx - \beta \max_s \{q(s)\} = E_s \int_0^{Q} \max\{0, p(x,s) - mc\}dx - \beta\, Q$$

Thus, the condition that characterizes the socially efficient capacity is

$$0 = \frac{dW}{dQ} = E_s \max\{0, p(Q,s) - mc\} - \beta \ ,$$

which is precisely the quantity supplied by a competitive industry.

Probably the main thing to understand about the competitive peak load environment is that efficiency and competition entail a positive probability of binding capacity. This is important to understand because regulators and reporters frequently appear to misunderstand that occasionally binding capacity is a natural part of a well-functioning marketplace, instead attributing binding capacity to monopoly, collusion or regulatory inefficiency.

4.1: Two-Period Peak Load Pricing

The basic peak-load pricing problem, pioneered by Marcel Boiteaux, considers two periods. The firm's profits are given by

$$\pi = p_1 q_1 + p_2 q_2 - \beta \max\{q_1, q_2\} - mc(q_1 + q_2)\,.$$

Prices equal to marginal costs are not sustainable, because a firm selling with price equal to marginal cost would not earn a return on the capacity, and thus would lose money and go out of business. A capacity charge is necessary. The question of peak load pricing is where the capacity charge should be allocated.

Demands are ordinarily assumed independent, but this is neither a good assumption nor a necessary one. Our previous analysis suggests how the solution will change, however, and so I will stick with independent demands for simplicity.

Social welfare is

$$W = \int_0^{q_1} p_1(x)dx + \int_0^{q_2} p_2(x)dx - \beta \max\{q_1, q_2\} - mc(q_1 + q_2).$$

The Ramsey problem is to maximize W subject to a profit condition. As always, write the Lagrangian

$$L = W + \lambda\, \pi.$$

Therefore,

$$0 = \frac{\partial L}{\partial q_1} = p_1(q_q) - \beta 1_{q_1 \geq q_2} - mc + \lambda\left(p_1(q_q) + q_1 p_1{}'(q_1) - \beta 1_{q_1 \geq q_2} - mc\right)$$

Or,

$$\frac{p_1(q_1) - \beta 1_{q_1 \geq q_2} - mc}{p_1} = \frac{\lambda}{\lambda + 1} \frac{1}{\varepsilon_1}$$

where $1_{q_1 \geq q_2}$ is the characteristic function of the event $q_1 \geq q_2$.

Similarly,

$$\frac{p_2(q_2) - \beta 1_{q_1 \leq q_2} - mc}{p_2} = \frac{\lambda}{\lambda + 1} \frac{1}{\varepsilon_2}$$

Note as before that $\lambda \to \infty$ yields the monopoly solution.

There are two potential types of solution. Let the demand for good 1 exceed the demand for good 2. Either $q_1 > q_2$, or the two are equal.

Case 1: $q_1 > q_2$.

$$\frac{p_1(q_1) - \beta - mc}{p_1} = \frac{\lambda}{\lambda + 1} \frac{1}{\varepsilon_1} \text{ and } \frac{p_2(q_2) - mc}{p_2} = \frac{\lambda}{\lambda + 1} \frac{1}{\varepsilon_2}.$$

In case 1, with all of the capacity charge allocated to good 1, quantity for good 1 still exceeds quantity for good 2. Thus, the peak period for good 1 is an extreme peak. In contrast, case 2 arises when assigning the capacity charge to good 1 would reverse the peak – assigning all of the capacity charge to good 1 would make period 2 the peak.

Case 2: $q_1 = q_2$.

An increase in output in either market requires a capacity increase, while a decrease in either market does not. The first order conditions become inequalities, of the form

$$0 \leq \frac{p_1(q_1) - mc}{p_1} - \frac{\lambda}{\lambda + 1} \frac{1}{\varepsilon_1} \leq \frac{\beta}{p_1} \text{ and } 0 \leq \frac{p_2(q_2) - mc}{p_2} - \frac{\lambda}{\lambda + 1} \frac{1}{\varepsilon_2} \leq \frac{\beta}{p_2}.$$

These must solve at $q_1 = q_2 = q$. The profit equation can be written

$p_1(q) - mc + p_2(q) - mc = \beta$

This equation shows that the capacity charge is shared across the two markets proportional to the inverse demand.

This theory represents two substantial simplifications: two periods and predictable demands. While stochastic demands present substantial complexities, handling more than two periods is straightforward and is considered next.

4.2: Multiperiod Peakload Pricing

Suppose there are n markets, and demand is given by $x_i(\mathbf{p})$ in market i where $\mathbf{p}=(p_1,\ldots,p_n)$. The peak load pricing problem is generally identified by a cost function $c(\mathbf{x}\bullet\mathbf{1}) - \beta \max \{x_i(\mathbf{p})\}$, but we will consider a general cost function $c(\mathbf{x})$. Profits are given by

$$\pi = \sum_{i=1}^{n} p_i x_i(\mathbf{p}) - c(\mathbf{x}) = \mathbf{p} \bullet \mathbf{x} - c(\mathbf{x}) .$$

As before, define the cross-price elasticity of substitution

$$\varepsilon_{ij} = \frac{p_j}{x_i}\frac{dx_i}{dp_j} ,$$

and let $\mathbf{E}$ be the matrix of elasticities. We maintain the assumption that demand $\mathbf{x}$ arises from a utility function and thus has a symmetric derivative. The regulatory pricing problem, which as before will subsume both the monopoly pricing problem and the socially efficient solution, is to

Consider the problem

$$\max u(\mathbf{x})\text{- } c(\mathbf{x}) \quad \text{s.t. } \mathbf{p}\bullet\mathbf{x} - c(\mathbf{x}) \geq \pi_0.$$

The symbol $\bullet$ is the standard Euclidean dot product. This formulation permits average costs to be decreasing. Write the Lagrangian

$$\Lambda = u(\mathbf{x}) - c(\mathbf{x}) + \lambda(\mathbf{p} \bullet \mathbf{x} - c(\mathbf{x})) = u(\mathbf{x}) - \mathbf{p} \bullet \mathbf{x} + (1+\lambda)(\mathbf{p} \bullet \mathbf{x} - c(\mathbf{x}))$$

Denote partial derivatives by subscripts: $c_i(\mathbf{x}) = \dfrac{\partial c(\mathbf{x})}{\partial x_i}$, and let $L_i = \dfrac{p_i - c_i}{p_i}$.

The lagrangian term λ has the interpretation that it is the marginal increase in welfare associated with a decrease in firm profit. Using Roy's identity,

$$0 = \frac{\partial \Lambda}{\partial p_i} = \lambda x_i + (1+\lambda)\sum_{j=1}^{n}(p_j - c_j)\frac{\partial x_j}{\partial p_i} = \lambda x_i + (1+\lambda)\sum_{j=1}^{n}(p_j - c_j)\frac{\partial x_i}{\partial p_j}$$

$$= \lambda x_i + (1+\lambda) x_i \sum_{j=1}^{n} \frac{p_j - c_j}{p_j} \varepsilon_{ij} .$$

Writing the first order conditions in vector form, we obtain the same general Ramsey price solution:

$$\mathbf{L} = -\frac{\lambda}{\lambda + 1}\mathbf{E}^{-1}\mathbf{1}.$$

Two complexities are added through the fact that *c* embodies peak-load pricing. First, there may be a kink in costs, and hence a discontinuity in marginal costs, around any set of prices in which two periods have the same quantity that is maximal over all periods. This problem is readily handled by precisely the same means as the two period case, using inequalities. Second, there could be several periods, all of which are peak periods and hence share the capacity cost among them. A computational strategy for addressing the second complexity is to assign all of the capacity charge to the highest demand period; if this produces a reversal (quantity below the peak demand of other periods), force the highest and the second highest to have the same quantity and assign the entire capacity charge to both of these periods. If the third highest now becomes the peak because the two highest have a larger marginal cost and hence a larger price, force the top three periods to have the same quantity, and so on.

The problem of random peaks is a much more challenging problem than the problem of multiple periods with predictable demand. There are several distinct problems in this category. First, unpredictable demand could mean demand that varies with the state of nature, as arose in the competitive model that began this section. For example, demand for electricity is higher on hot days as people turn on their air conditioners. In contrast, the demand for airline seats fluctuates not only because individuals may demand more travel but also because the set of individuals demanding seats evolves over time in a random fashion. The latter problem is much more challenging than the former because a firm cannot contract with the potential buyers in advance. When demand fluctuates but the set of consumers is identifiable, a firm might contract in advance, determining the quantity each consumer gets as a function of the state of nature that arises.

4.3: Priority Pricing

The peak load problem is essentially a cost allocation problem. It has an efficiency aspect, in that pricing matters to relative demand, but that efficiency aspect is incorporated in a familiar way, using inverse elasticities. The priority pricing problem introduced by Robert Wilson has a superficial similarity to the peak load problem – when capacity is reached, who should be rationed? Implicitly, the peak load formulation implies a spot market, so that each market is rationed efficiently. In many circumstances, it is not possible to use prices ex post to ration the market. For example, absent smart appliances, it is difficult for homeowners to adjust electric demand in real time as prices vary – homeowners aren't even informed about the abrupt price changes. Priority pricing is a means of contracting in advance when capacity, or demand, is stochastic.

At this time, the problem of stochastic demand and priority pricing has not been adequately addressed. In particular, with stochastic demand, there is an issue of whether all customers are able to participate in the ex ante priority market.

Consider a case of a continuum of consumers, each of whom desires one unit. As will become clear, it doesn't matter if some consumers desire multiple units – each unit can be treated as if demanded by a separate consumer. Rank the consumers by their valuations for the good, so that the q^{th} consumer has a value $p(q)$ for the good, and p is downward sloping.

The quantity available is a random variable with distribution F. Priority pricing is a charge schedule c which provides a unit to a customer paying $c(q)$ whenever realized supply is q or greater.

It is a straightforward exercise to calculate the incentive compatible c schedule. A customer of type q should choose to pay $c(q)$ for the q^{th} spot in the priority list. This leads to the incentive constraint:

$$u(q) = (p(q) - c(q))(1 - F(q)) \geq (p(q) - c(\hat{q}))(1 - F(\hat{q})).$$

The envelope theorem gives

$$u'(q) = p'(q)(1 - F(q)).$$

It is a straightforward exercise to demonstrate that the first order condition is sufficient; see 3.2:. Let $F(H)$=1, so that $u(H)$=0. Then

$$(p(q) - c(q))(1 - F(q)) = u(q) = -\int_q^H u'(s)ds = -\int_q^H p'(s)(1 - F(s))ds = p(q)(1 - F(q)) - \int_q^H p(s)f(s)ds$$

Thus,

$$c(q) = \int_q^H p(s)\frac{f(s)}{1 - F(q)}ds = E[spot\ \ price \mid p(s) \geq p(q)].$$

Revenues to the firm from the priority pricing are

$$R = \int_0^H c(q)(1 - F(q))dq = \int_0^H \int_q^H p(s) f(s)\, ds\, dq = \int_0^H qp(q) f(q)\, dq.$$

This is the revenue associated with a competitive supply; a monopolist might have an incentive to withhold capacity to boost prices. How does a monopolist do so? Withholding of capacity has the property of changing the distribution of available supply, in a first order stochastic dominant manner. In particular, the monopolist can offer any distribution of capacity G, provided G≥F. What is the monopolist's solution? Rewrite R to obtain

$$R = \int_0^H qp(q)g(q)dq = \int_0^H MR(q)(1-G(q))dq.$$

Provided marginal revenue MR is single-peaked,

$$G = \begin{cases} F \ if \ \ MR \geq 0 \\ \\ 1 \ if \ \ MR < 0 \end{cases}.$$

That is, the monopolist cuts off the capacity at the monopoly supply, and otherwise supplies the full amount.

4.4: Matching Problems

Priority pricing is a solution to a matching problem, matching the high value buyers with capacity. Many other problems have this feature, that it is desirable to match high types with high types and low types with low types. Such models have been used as models of marriage, employment, university admissions, incentive contracts, and other categories. Wilson examines not just the continuum matching, in which each probability of service interruption is separately priced, but also finite groups. Rather than offer a continuum of categories, consider offering just two – high priority service and low priority service. How well does such a priority service do?

The answer is: remarkably well. Consider first the linear demand case with a uniform distribution of outages. Perfect matching gets a payoff

$$\int_0^1 p(q)(1-q)dq = \int_0^1 (1-q)^2 dq = \frac{1}{3}.$$

No matching – that is a random assignment – produces an expected value of ¼, a fact which is evident from

$$\int_0^1 p(q)dq \int_0^1 (1-q)dq = \left(\int_0^1 (1-q)dq \right)^2 = \frac{1}{4}.$$

Now consider two groups of equal size. The high value group has an average value of ¾, and is served with probability $\int_0^{1/2} 2q dq + \int_{1/2}^1 1 dq =$ ¾. The low value group has average value ¼ and is served with probability ¼. Thus, the expected value from two categories is

$\frac{1}{2}\left(\frac{9}{16}+\frac{1}{16}\right)=\frac{5}{16}$. Note that 5/16 is 75% of the way from ¼ to 1/3! That is, a single group captures 75% of net value of a continuum of types.

The linear payoff/uniform distribution is special; however, McAfee (2002) shows that, provided a common hazard rate assumption is satisfied, two groups of distinguished by being above or below the mean generally captures 50% or more of the possible gains over no priority pricing. That is, using two classes is sufficient to capture a majority of the gains arising from priority pricing. More generally, Wilson shows that the losses from finite classes are on the order of $1/n^2$, although this doesn't translate into a specific bound for any given n.

4.5: Dynamic Pricing

The problem of dynamic pricing, also known as yield management and revenue management, is to adjust the price or prices of a good in limited supply with imperfect information about the realization of demand. Telluri and van Ryzin (2004) is a very eloquent and substantial book that provides an in depth introduction to the topic. Here, we consider a special theory, developed by Gallego and van Ryzin (1994) and with some more economic results found in McAfee and te Velde (2006).

A seller has a fixed stock of a perishable product to sell. Time is assumed to start at zero and end at T. If not sold, the product perishes at T, which might occur because an airplane departs – it is challenging to board in-flight – or the product for sale is a rental on a particular date, like a hotel room. For simplicity, discounting is assumed away; this is the right assumption when payment is made at T, as is common in hotel rooms, but incorrect when payment is made upon purchase, as with airplane seats. The marginal cost of the product is assumed constant at c. This value should be interpreted to include not just the actual marginal costs of service like cleaning a hotel room or peanuts served on a flight, but also the lost business – a customer who flies on a particular flight won't fly on an alternative flight offered by the same airline. Potential customers are assumed to demand a single unit, and their willingness to pay is given by a cumulative distribution function F. Let $\lambda(t)$ be the arrival probability of customers per unit of time, assumed constant. Time will start at 0 and end at T.

The value of having n items for sale at time t is denoted by $v_n(t)$. Clearly having nothing to sell conveys zero value. Moreover, if not sold by T, any inventory of items also has zero value. Together, these imply:

$$v_0(t) = v_n(T) = 0.$$

To compute the equation governing v, we employ a simple trick. Consider a small increment of time, Δ, beyond a given time t. Given a flow rate of customer arrivals $\lambda(t)$, the probability of a customer arriving is $\lambda(t)\ \Delta$. With probability about $1-\lambda(t)\Delta$, no customer arrives, so that the current value becomes $v_n(t+\Delta)$. Alternatively, with probability $\lambda(t)\Delta$, a customer arrives and the firm either makes a sale or does not. For price p, the sale occurs with probability $1-F(p)$. When a sale occurs, the profit is becomes $p - c + v_{n-1}(t+\Delta)$, because the inventory is decreased by one. Summarizing, and suppressing the dependence of λ on t:

$$v_n(t) = \max_p \ (1-\lambda\Delta)v_n(t+\Delta) + \lambda\Delta\big((1-F(p))(p-c+v_{n-1}(t+\Delta)) + F(p)v_n(t+\Delta)\big)$$

or

$$v_n(t) - v_n(t+\lambda\Delta) = \lambda\Delta \max_p \ (1-F(p))(p-c+v_{n-1}(t+\Delta)) - v_n(t+\Delta))$$

Therefore, dividing by Δ and sending Δ to zero,

$$-v_n'(t) = \lambda(t)\max_p \ \ (1-F(p))(p-c+v_{n-1}(t)-v_n(t))$$

The expression for $v_n'(t)$ is composed of two terms. First, there are profits from a sale, *p*–*c*. Second, there is the lost option of selling the unit in the future, an option that has value $v_n(t)-v_{n-1}(t)$. With a convenient choice of distribution *F*, it is possible to provide an explicit solution. Let

$$F(p) = 1-e^{-ap}.$$

Note that the function (1-e^{-ap})(*p*–*mc*) is maximized at $p^* = {}^1\!/_a + mc.$[5] Therefore, the profit-maximizing price is

$$p_n^*(t) = \frac{1}{a} + c + v_n(t) - v_{n-1}(t),$$

and his can be plugged into the expression for $v_n'(t)$ to obtain

$$v_n'(t) = -\lambda(t)e^{-a\left(\frac{1}{a}+c+v_n(t)-v_{n-1}(t)\right)}\frac{1}{a}.$$

This equation for $v_n'(t)$ is solvable by induction; it is easiest to display the solution and verify it. To do so, it is useful to introduce notation

$$\beta(t) = e^{-1-ac}\int_t^T \lambda(s)ds\,.$$

[5] Let *h*(*p*)= (1-e^{-ap})(*p*–*mc*). $h'(p) = e^{-ap}(1-a(p-c))$ $h'(p)=0$ implies $h''(p) \le 0$. Thus, every extreme point is a maximum, and so if there is an extreme point, it is the global maximum. Moreover, *p** is an extreme point.

The function β represents the expected number of future customers willing to pay the static monopoly price ${}^{1}\!/_{a}+c$. At time t, the expected number of buyers willing to pay the monopoly price is $\beta(t)$. The second piece of notation is

$$B_n(t) = \sum_{j=0}^{n} \frac{\beta(t)^j}{j!}.$$

Gallego and van Ryzin (1994) show that $v_n(t) = \frac{1}{a} Log\left(B_n(t)\right)$. This is readily proved by checking that $v_n(T)=0$ (which holds because $B_n(T)=1$) and that the differential equation for $v_n'(t)$ is satisfied. In addition, the equation for price is readily computed from $p_n^*(t) = \frac{1}{a} + c + v_n(t) - v_{n-1}(t)$, and has the useful expression

$$e^{-ap_n^*} = e^{-1-ac}\,\frac{B_{n-1}(t)}{B_n(t)}.$$

Suppose there is an initial capacity K. Let $q_k(t)$ be the probability that there are k units left in inventory for sale at time t. The value of q_k evolves according to inflow (sales from the inventory k+1) and outflow (sales from the inventory k). Thus,

$$\begin{aligned}
q_k'(t) &= \lambda(t)(1-F(p_{k+1}(t))q_{k+1}(t) - \lambda(t)(1-F(p_k(t))q_k(t) \\
&= \lambda(t)e^{-ap_{k+1}(t)}q_{k+1}(t) - \lambda(t)e^{-ap_k(t)}q_k(t) \\
&= \lambda(t)\frac{e^{-1-ac}B_k(t)}{B_{k+1}(t)}q_{k+1}(t) - \lambda(t)\frac{e^{-1-ac}B_{k-1}(t)}{B_k(t)}q_k(t) \\
&= -\frac{B_{k+1}'(t)}{B_{k+1}(t)}q_{k+1}(t) + \frac{B_k'(t)}{B_k(t)}q_k(t)
\end{aligned}$$

Given a capacity K at time 0, $q_K(0)=1$, $q_{K+1}(t)=0$, and thus $q_K(t) = \frac{B_K(t)}{B_K(0)}$.

This is used as the base of an induction to establish $q_n(t) = \frac{(\beta(0)-\beta(t))^{K-n}B_n(t)}{(K-n)!B_K(0)}$. The details of the induction are presented in Appendix 1.

The inventory or expected number of unsold items, $K - n$, satisfies

$$E(K-n)=\sum_{n=0}^{K}(K-n)q_n(t)=\sum_{n=0}^{K-1}(K-n)\frac{(\beta(0)-\beta(t))^{K-n}B_n(t)}{(K-n)!\,B_K(0)}$$

$$=\sum_{n=0}^{K-1}\frac{(\beta(0)-\beta(t))^{K-n}B_n(t)}{(K-1-n)!\,B_K(0)}$$

$$=\frac{(\beta(0)-\beta(t))B_{K-1}(0)}{B_K(0)}\sum_{n=0}^{K-1}\frac{(\beta(0)-\beta(t))^{K-1-n}B_n(t)}{(K-1-n)!\,B_{K-1}(0)}$$

$$=\frac{(\beta(0)-\beta(t))B_{K-1}(0)}{B_K(0)}.$$

This model is a monopoly pricing model. How does it compare to the efficient solution? An efficient solution in this model has the property that the value function maximizes the gains from trade rather than the profit. The value function, then, satisfies

$$S_n(t)=\max_p\ (1-\lambda\Delta)S_n(t+\lambda\Delta)+\lambda\Delta\big((1-F(p))(G(p)+v_{n-1}(t+\lambda\Delta))+F(p)v_n(t+\lambda\Delta)\big),$$

and thus

$$-S_n'(t)=\lambda(t)\max_p\ \ (1-F(p))(G(p)+S_{n-1}(t)-S_n(t))$$

where G(p) is the consumer surplus, plus seller profit, conditional on the consumer's value exceeding p. Given the demand specification,

$$(1-F(p))G(p)=\int_p^{\infty}(1-F(x))dx+(p-c)(1-F(p))$$

$$=\int_p^{\infty}e^{-ax}dx+(p-c)e^{-ap}=(p+{}^{1}\!/_{a}-c)e^{-ap}\text{, so that } G(p)=p+{}^{1}\!/_{a}-c\,.$$

Thus the efficient solution is the solution a monopoly whose costs were reduced by 1/a, the static monopoly profit, would choose. That is, for any monopoly solution, there is another demand function under which that monopoly solution would be efficient! Consequently, the price variation in the solution should be attributed to efficient rationing, rather than monopoly pricing; monopoly pricing in this model only has the effect of shifting the price up by the constant 1/*a*.

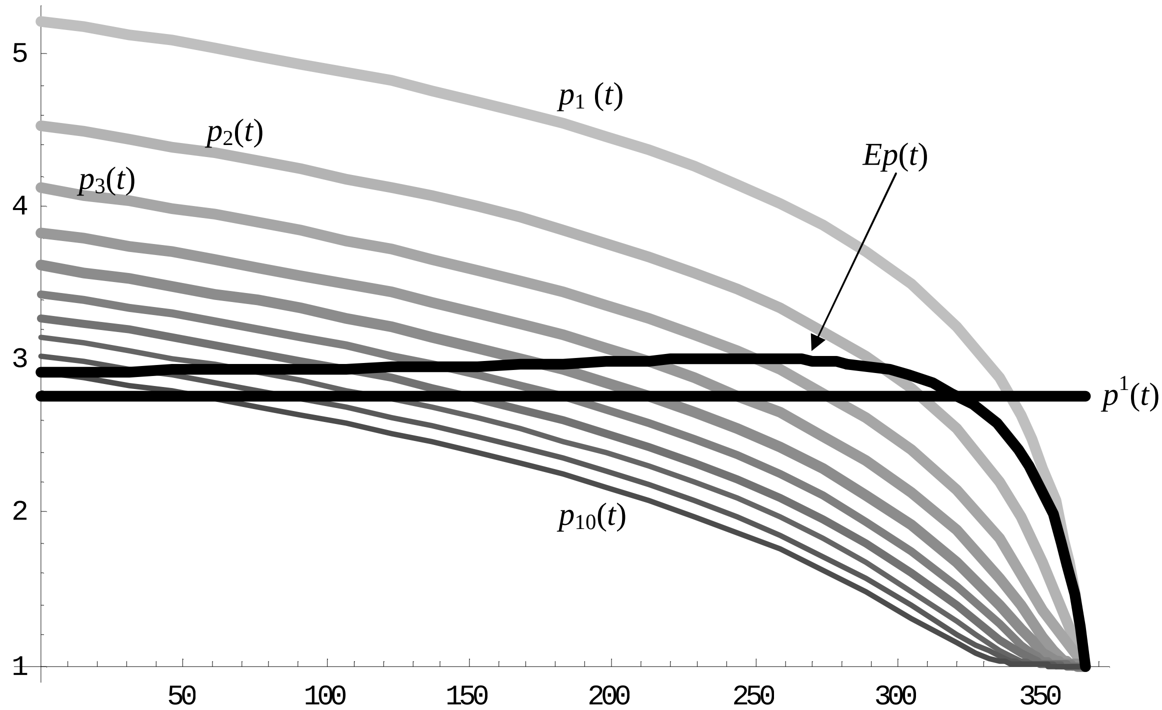

Figure 6: Selling 10 seats

Figure 6 illustrates the solution when λ is a constant, and K=10. The different price paths $p_n(t)$ are illustrated in various grey lines, some of which are named. In addition, the expected price conditional on availability is denoted by Ep, and the price that a firm that had to set a fixed price and not change it (but isn't required to sell beyond capacity) is denoted by p^1. The interesting thing here is that the expected price isn't much different than p^1 for most of the time; the theorem is that for any demand, as the time and number of items diverge proportionally, the price path converges to a constant. Thus, most of the gains from dynamic pricing arise from the last few items; in selling hundreds of hotel rooms or airplane seats for a particular night, dynamic pricing has a modest effect.

5: Price Dispersion

George Stigler's 1961 "The Economics of Information" begins with:

> *One should hardly have to tell academicians that information is a valuable resource: knowledge is power. And yet it occupies a slum dwelling in the town of economics.*

Stigler documents the existence of price dispersion empirically, and proceeds to sketch various theories that might be applicable to price dispersion. Stigler is careful to distinguish price dispersion, which results from imperfect information on the part of consumers, from price discrimination, which involves distinct prices for different types of consumer preferences.

There are various reasons that price dispersion might arise. Clearly some consumers must be willing to pay higher prices than others, that is, distinct consumers have distinct reservation prices. (Note that a reserve price – the minimum bid in an auction – is a very different concept from a reservation price.) Differences in reservation prices can arise because of differences in knowledge – some consumers have access to price lists, others don't – or search costs. If some consumers have lower search costs, then these consumers will search for lower prices.

Models based on search costs usually require differences in the firms as well as differences in search costs on the part of buyers. The reason is that, with differences in consumer search, identical firms will tend to respond in an identical fashion to the generated demand, and the price dispersion is degenerate. An extreme example of this phenomenon is the Diamond Paradox. (Diamond, 1971.) Suppose that there is global minimum to search costs γ, that is, all consumers have search costs in excess of $\gamma>0$. Search costs arise per store – each store sampled costs γ or more. One must incur the search cost to find out the price in this model, as opposed to a cost of visiting a store with a known price. Then we conclude that all firms charge the monopoly price. The proof is straightforward. Let L be the lower bound on the distribution of prices and suppose $L < M$, the monopoly price. If a firm charges $\text{Min}\{L + \tfrac{1}{2}\gamma, M\}$, it is strictly better than charging any lower price, because no consumer will reject the price $L + \tfrac{1}{2}\gamma$ that will accept L, because the cost of obtaining an additional price, even if it is certain to be L, is at least γ! Thus, even if all consumers have a very low but positive cost of search, the equilibrium involves all firms charging the monopoly price. What is paradoxical about this result is that the equilibrium is discontinuous in the search cost – the equilibria with zero search costs (price equals marginal cost) and any positive search costs are dramatically different.

The literature in the 1970s focused on models with varying search and production costs as a means of generating price dispersion. (See Carlson and McAfee (1983)for an example and a list of references.) The more modern and economical approach involves generating price dispersion from identical firms via randomization. This approach is intellectually more satisfying because we don't need a story for why the firms have different costs; moreover, applications involving entry of identical firms become possible. Finally, it is possible to generate the consumer informational asymmetries endogenously, as Varian (AER, 1980) does.

5.1: The Butters Model

Suppose n firms send advertisements to a proportion α of the population with price offers and these offers are distributed randomly. Each consumer chooses the price offer that is lowest. It will not be an equilibrium for each firm to send out the same price offer – it would pay either firm to undercut it by a fraction of a penny. So we will look for a mixed strategy – each firm sends out a random price offer. Let $F(p)$ be the probability that a price offer is not more than p. Given price p, the quantity purchased by the consumer is $q(p)$. Let $R(p) = (p - c)q(p)$ be the profit per consumer sold, and M the smallest maximizer of R. It is useful to assume R is increasing for $p<M$ and that assumption is maintained here.

A firm's profits per consumer are

$$\pi(p) = R(p)(1-\alpha + \alpha(1-F(p)))^{n-1}$$

since it beats another firm if the consumer doesn't receive an offer from that firm (1-α) or if the offer received from that firm has a price in excess of p. Each firm must earn the same profits for every price in the support of the distribution of prices, for otherwise firms would never choose the low profit prices. The monopoly price M must be in the limit of the support, because if the maximum offer were strictly less than that, it would pay to offer M [the highest offer is only accepted if no other offer is received; it is best to make the most money possible in this event]. Finally, no firm will ever offer more than M because an increase in price beyond M does not increase revenue conditional and lowers the probability of selling. Thus,

$$\pi(p) = \pi(\text{M}) = R(M)\,(1-\alpha)^{n-1}$$

Therefore,

$$1-\alpha F(p) = 1-\alpha+\alpha(1-F(p)) = (1-\alpha)\left(\frac{R(M)}{R(p)}\right)^{\frac{1}{n-1}}.$$

or

$$F(p) = \frac{1}{\alpha}\left[1-(1-\alpha)\left(\frac{R(M)}{R(p)}\right)^{\frac{1}{n-1}}\right].$$

The lowest price, L, is the price such that $F(L)=0$, or $R(L) = (1-\alpha)^{n-1}R(M)$.

The distribution of prices is Prob(offer$\leq p$) = $1-[1-\alpha+\alpha(1-F(p)]^n$. The expected profits in the model are

$$E\pi = \int_L^M R(p)n(1-\alpha F(p))^{n-1}\alpha f(p)dp = \int_{p_0}^{v} R(M)(1-\alpha)^{n-1} n\alpha f(p)dp = n\pi\alpha = R(M)(1-\alpha)^{n-1} n\alpha.$$

In this model, some consumers receive no offers, and no profits are made on those consumers. Low profits for this reason is not an indication of efficiency, indeed, quite the contrary. Consequently, a measure of performance of the industry is the profits per customer served, and the probability a customer is served, which is equivalent to them receiving at least one price offer, is $1-(1-\alpha)^n$. Thus, expected profits per customer served

$$\frac{E\pi}{1-(1-\alpha)^n} = R(M)\frac{(1-\alpha)^{n-1}n\alpha.}{1-(1-\alpha)^n}$$

This is interpretable in an alternative, sensible way: the measure of market profits is the level of monopoly profits $R(M)$, times the probability of receiving exactly one offer, conditional on receiving at least one offer. That is, the same measure of market performance would arise when firms could condition their offers on whether consumers received multiple offers or not, in which case the firms would charge M when the customer received no competing offers, and c otherwise.

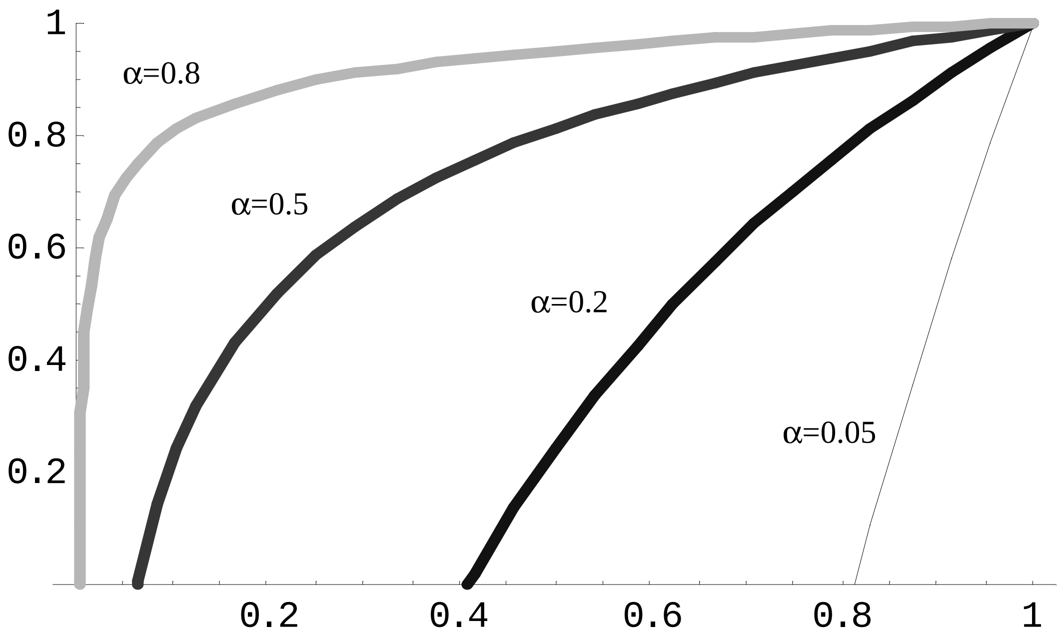

Figure 7: The Probability Distribution for Selected Values of α, with *c*=0 and *n*=5.

This measure of market profits goes from one at $\alpha=0$ to zero at $\alpha=1$. It is a decreasing function of α, so more searchers reduces profits per served customer, as one would expect. Profits per customer are decreasing in n, the number of firms. The proof that it is increasing in n has a trick. Define $x=(1-\alpha)^n$. Then $n = \frac{\log(x)}{\log(1-a)}$, which is a decreasing function of x. Note that

$\frac{(1-\alpha)^{n-1} n\alpha}{1-(1-\alpha)^n} = \frac{x\log(x)}{1-x} \frac{\alpha}{(1-\alpha)\log(1-\alpha)}$. It is readily established that $\frac{x\log(x)}{1-x}$ is decreasing and negative in x for $0<x<1$, which makes the product increasing (recall $\log(1-\alpha) < 0$) in x, and therefore decreasing in n.

Figures 6 and 7 show the distribution of prices and the distribution of best prices for a variety of values of α, in the special case where q is either 1 if price is below M, or zero otherwise.

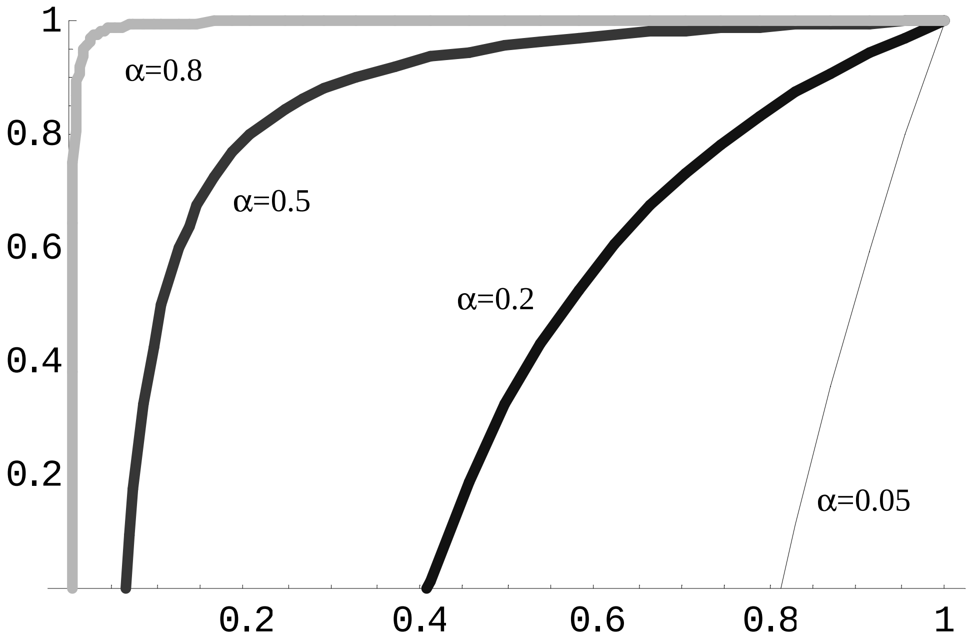

Figure 8: The Distribution of the Best Price, for the Same Values of α.

To look at merger in the Butters model, we would have to consider asymmetric α - a much more difficult problem. However, the Butters model produces applicable imperfect price competition in some circumstances, and also produces random prices - a model of sales.

Note that advertising in the Butters model is informative, telling consumers where they can get a low price.

The solution to the Butters model is equivalent to the auction model in which bidders have independently distributed values in {0,1}, and α is the probability that a given bidder has value equal to 1.

The Butters model can be varied in a straightforward way by letting buyers either know one price (chosen randomly) or all the prices, a model which is developed in Varian's AER paper. In this case, suppose β is the probability of observing only one price. These "loyal" customers are evenly distributed across the n firms; the other customers, "shoppers," buy only from the least expensive firm. Then firms face a profit function:

$$\pi(p) = R(p)\left(\beta/_{n} + (1-\beta)(1-F(p))^{n-1}\right).$$

As before, $\pi(p) = \pi(M)$ and thus

$$R(p)\left(\beta/_{n} + (1-\beta)(1-F(p))^{n-1}\right) = \frac{R(M)\beta}{n}.$$

This solves for F in a straightforward way.

Figure 9: Some consumers are willing to pay a higher price, while some shop for the lower price.

Varian goes on to endogenize the number of consumers who choose to be informed given a cost of being informed, an important development because it makes the entire model self-contained and rational. Since shoppers get a better price than loyal customers, β is determined by indifference between paying a cost to learn all the prices and choosing a store at random.

As posited, the price dispersed models are beautiful models that are not useful for industrial organization applications. To make them more useful, it is necessary to endogenize the access of the firms to consumers, that is, endogenize the value of α, and that requires handling asymmetries. In McAfee (1992) an asymmetric model is developed. In this model, firms have availability rate α_i, ranked from largest to smallest, so that $\alpha_1 \geq \alpha_2 \geq \ldots$

With a substantial amount of work, one can show that there is an equilibrium, that firm 1 has a mass point at M if $\alpha_1 > \alpha_2$, that the firms with lower availability randomize over intervals with lower prices, and finally that profits per unit of availability are the same for all firms, but that the largest firm enters asymmetrically into the profit equation. For any firm i,

$$\pi_i = \alpha_i R(M) \prod_{j \neq 1} (1-\alpha_j).$$

Now introduce a cost $c(\alpha)$ of availability. Given the multiplicative nature of probabilities, the size of a scale economy is given by the cost saving associated with combining the operations of two entities, which is $c(\alpha)+c(\beta)-c(1-(1-\alpha)(1-\beta))$. Consequently, availability has increasing returns to scale whenever $c(\alpha)+c(\beta)-c(1-(1-\alpha)(1-\beta))$ is increasing in α or β. It turns out that increasing returns to scale are equivalent $(1-\alpha)c'(\alpha)$ being decreasing in α. Constant returns to scale involve $(1-\alpha)c'(\alpha)$ being constant, which implies $c(\alpha)=-\theta\log(\alpha)$.

It turns out that there is a pure strategy equilibrium, which involves $\alpha_1>\alpha_2=\alpha_3=\ldots=\alpha_n$. If there are increasing returns to scale, $\alpha_1>2\alpha_2$. If there is a diseconomy of scale, $\alpha_1<2\alpha_2$. Thus, an initially symmetric model generates an asymmetric equilibrium, and no symmetric pure strategy equilibria.

The closed form solution to the price dispersed equilibrium makes it a natural vehicle for applications in industrial organization theory where price competition, rather than quantity or capacity competition, is the natural assumption.

5.2: Search

How should a consumer searching for a low price go about it? How should some seek to find a job? How should someone bid on eBay when many sellers are offering the same or similar items in auctions that have distinct ending times? These questions are the province of search theory.

Fom the perspective of a potential buyer, prices offered at any one store or at any given moment are random variables. For simplicity assume the offered price has a probability density function $f(p)$. The theory assumes that buyers incur a search cost to obtain a price quote, which might arise if the buyers has to visit a store in person, telephonically or virtually. The cost includes your time and any other costs necessary to obtain a price quote. In this instance, the consumer will set a *reservation price*, which is a maximum price they will pay without visiting another store. That is, if a store offers a price below p^*, the consumer will buy, and otherwise they will visit another store, hoping for a better price. The optimality of such a rule follows from the fact that the further search has an expected optimal value; the buyer is deciding between buying now and further search. A lower price today obviously favors buying now.

Suppose that the cost of search is c. Let p^* denote the reservation price and $J(x)$ represent the expected total cost of purchase (including search costs) if a reservation price of x is used. Then J must equal

$$J(x)=c+\int_0^x pf(p)dp+\int_x^\infty J(x)f(p)dp\,.$$

This equation arises because the current draw (which costs c) could either result in a price less than x, in which case observed price, with density f, will determine the price paid p, or the price will be too high, in which case the consumer is going to take another draw, at cost c, and on average get the average price $J(x)$. It is useful to introduce the cumulative distribution function F, with $F(x) = \int_0^x f(p)dp$. Note that something has to happen, so $F(\infty)=1$.

We can solve the equality for $J(x)$, $J(x) = \dfrac{\int_0^x pf(p)dp + c}{F(x)}$.

The function J has a simple interpretation. The expected price $J(x)$ is composed of the expected price, $\int_0^x p\dfrac{f(p)}{F(x)}dp$, which is the average price conditional on that price being less than x, plus a term depending on search costs. Note that $\dfrac{f(p)}{F(x)}$ is the density of the price conditioned on the price being less than x. The second term is $\dfrac{c}{F(x)}$ is the expected search costs, and it arises because $\dfrac{1}{F(x)}$ is the expected number of searches. To see this, note that the search takes n trials with probability $F(x)(1-F(x))^{n-1}$. Thus the expected number of searches is $\sum_{n=1}^{\infty} nF(x)(1-F(x))^{n-1} = \dfrac{1}{F(x)}$. (It is an exercise to prove the equality.)

But what reservation price of x minimizes cost of purchase $J(x)$?

$$J'(x) = x\frac{f(x)}{F(x)} - \frac{f(x)\int_0^x pf(p)dp + c}{F(x)^2}$$

$$= \frac{f(x)}{F(x)}\left(x - \frac{\int_0^x pf(p)dp + c}{F(x)} \right) = \frac{f(x)}{F(x)}\left(x - J(x)\right).$$

Thus, if $x<J(x)$, J is decreasing, and it lowers cost to increase x. Similarly, if $x>J(x)$, J is increasing in x, and it reduces cost to decrease x. Thus, minimization occurs at the reservation price p^* satisfying $p^*=J(p^*)$.

Moreover, there is only one such solution to the equation $p^*=J(p^*)$ in the range where f is positive. To see this, note that at any solution to the equation $p^*=J(p^*)$, $J'(p^*)=0$ and

$$J''(p^*) = \frac{d}{dp^*}\left(\frac{f(p^*)}{F(p^*)}\left(p^* - J(p^*)\right)\right)$$

$$= \left(\frac{d}{dp^*}\frac{f(p^*)}{F(p^*)}\right)(p^* - J(p^*)) + \frac{f(p^*)}{F(p^*)}(1 - J'(p^*)) = \frac{f(p^*)}{F(p^*)} > 0.$$

This equation shows that J takes a minimum at any solution to $p^*=J(p^*)$, since its first derivative is zero and its second derivative is positive. Moreover, because every solution is a minimum, the solution is unique; two minima must have a local maxima between them. Otherwise, J would have to be both increasing and decreasing on an interval between two consecutive minima, since J is increasing to the right of the first solution, and decreasing to the left of the second solution. Consequently, the equation $p^*=J(p^*)$ has a unique solution, that minimizes the cost of purchase.

Consumer search to minimize cost implies setting a reservation price equal to the expected total cost of purchasing the good, and purchasing whenever the price offered is lower than that level. That is, it is not sensible to "hold out" for a price lower than what you expect to pay on average, although of course such a holdout strategy might be well useful in a bargaining situation.

Example (Uniform): Suppose prices are uniformly distributed on the interval $[a,b]$. For p^* in this interval,

$$J(p^*) = \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)} = \frac{\int_a^{p^*} p\frac{dp}{b-a} + c}{\frac{p^*-a}{b-a}}$$

$$= \frac{½(p^{*2} - a^2) + c(b-a)}{p^*-a} = ½(p^*+a) + \frac{c(b-a)}{p^*-a}.$$

Since $p^*=J(p^*)$, $p^* = a + \sqrt{2c(b-a)}$.

As $c \rightarrow 0$, $p^* \rightarrow a$, that is, as the search costs go to zero, one holds out for the lowest possible price. This is reasonable in the model, but in real world situations, delay may cause discounting, which isn't accounted for in the model. In addition, $p^* < b$, the maximum price, if $2c<(b-a)$. Put another way, if the *most* you can save by a search is twice the search cost, then search is not optimal, because the expected gains from search will be half the maximum gains (thanks to the uniform distribution) rendering search unprofitable.

A property of the uniform distribution that generalizes to any distribuiton is that the expected price is a concave function of the cost of search. To see this, define a function

$$H(c) = \min_x J(x) = \min_x \frac{\int_0^x pf(p)dp + c}{F(x)}.$$

Since $J'(p^*) = 0$,

$$H'(c) = \frac{\partial}{\partial c} J(p^*) = \frac{1}{F(p^*)} > 0$$

Moreover, *p** is increasing in *c*, from which it follows that *H* is concave. This means that the effects of an increase in *c* are passed on at a decreasing rate. This follows from differentiating $J(p^*) - p^*=0$ to find

$$\frac{dp^*}{dc} = \frac{\partial(J(p^*) - p^*)/\partial c}{\partial(J(p^*) - p^*)/\partial p^*} = -\frac{1/F(p^*)}{-1} = 1/F(p^*).$$

It helps to note that $J'(p^*) = 0$ in this calculation. This receives the interpretation that a small cost increase has an effect on the reservation price as if the search strategy is unchanged (same *p**), so that that cost increase is applied to the existing expected number of searches $1/F(p^*)$.

6: Experience Goods

Some goods have to be tried to evaluate accurately. Even though the critics dislike a movie, you may love it, and conversely. Restaurant appeal is not universal. Shoes that fit you perfectly are uncomfortable to someone else. Such goods are called experience goods, meaning that you must try them to actually know if they are good or not.

A seller offering an experience good is subject to the beliefs of customers about the quality of her good. If customers believe the good to be of low quality, they will be unwilling to pay a high price for it; if no one buys the good, its intrinsic qualities are not discovered, thus perpetuating the belief independent of the actual quality. For the seller of a low quality good, this is as it should be, but for the seller of a high quality good, it represents a quandary. How can the seller of a high quality experience good communicate the quality of the good?

A seller seeking repeat business can offer an introductory discount as a means of communicating the quality of the good. For example, a restaurant may offer a discount for the first meal, which has the effect of subsidizing experimentation. Such a subsidy to experimentation will communicate quality because the seller can only recoup the subsidy through repeat business, which will not be forthcoming if the meal is in fact miserable.

Models of introductory pricing cleave into two categories: adverse selection and moral hazard. The adverse selection models represent the situation where the type of sellers varies – some are high quality and some are low quality. For example, there are good and bad attorneys and it is difficult to establish the quality of attorney by reading their win rate since every settlement is counted as a win by both sides, and most cases settle. Similarly, chefs vary in skill level. In contrast, the case of moral hazard arises when the seller decides, on a period by period basis, whether to sell the high quality good or a shoddy low quality imitation. For example, a restaurant might choose to use high quality, flavorful ingredients or rotten floor-sweepings in their soup. This would be a case of moral hazard because any restaurant faces such a choice. As the examples suggest, most real world settings probably involve a mix of both moral hazard and adverse selection.

6.1: Adverse Selection

Consider a product which consumers can purchase in every period. There are two quality levels, H and L, with unit costs $c_H > c_L$. Let willingness to pay be given by p_H for the high quality good and p_L for the low quality good. Assume that $p_H - c_H > 0 > p_L - c_L$, so that optimally, only high quality products are sold. If the low quality good is optimal, there will be no challenge for the marketplace. Once a manufacturer has built a plant, the quality of its output remains constant forever, but consumers are initially uninformed about which type of plant it built. The firm that builds the high quality plant gets repeat business, provided that consumers try the good. The low quality plant gets no repeat business, because at any price that covers costs, consumers are unwilling to purchase.

Suppose the manufacturer is going to post price p_0 initially and then post price p thereafter. Let δ be the discount factor applied to each future period's profits. The seller's profit is

$$p_0 - c_H + \sum_{t=0}^{\infty} \delta^t \ (p - c_H) = p_0 - c_H + \delta \frac{p - c_H}{1-\delta}$$

Imitating this strategy, a low quality manufacturer obtains $p_0 - c_L$. Thus, as long as there are prices p_0 and p that consumers are willing to pay so that

$$p_0 - c_L \leq 0 \leq \ p_0 - c_H + \delta \frac{p - c_H}{1-\delta}.$$

then an *introductory offer* sufficiently low will guarantee that every firm produces the high quality. In particular, if any prices work, then $p_0 = c_L$ and $p = c_H$ will work, since these weaken the constraints maximally. Thus, an introductory offer can guarantee quality when

$$0 \leq \ c_L - c_H + \delta \frac{p_H - c_H}{1-\delta},$$

or

$$c_H \leq \ (1-\delta)c_L + \delta p_H .$$

In particular, if the discount factor δ is close enough to one, this is automatically satisfied from the hypothesis that the high quality good is efficient.

6.2: Moral Hazard

The moral hazard decision differs from the adverse selection problem because the seller can choose at any time to lower the quality of the good. Suppose that again there are two quality levels, H and L, with unit costs $c_H > c_L$. Again let willingness to pay be given by p_H for the high quality good and p_L for the low quality good. Assume that $p_H - c_H > 0 > p_L - c_L$, so that optimally, only high quality products are sold. Again, it is no problem inducing a seller to offer the low cost good; the challenge for a marketplace is to induce a seller to offer the high quality good.

Suppose that buyers use a *grim trigger strategy*, which involves never buying from a seller that offered the low quality. This is about as severe a punishment as buyers can choose, and has a certain psychological plausibility to it. If the seller charges a price p, it pays to never offer the low quality if

$$p - c_L \leq p - c_H + \delta \frac{p - c_H}{1-\delta}.$$

This inequality says that offering the low quality good, but never selling again, produces lower profits than offering the high quality good and gaining repeat business. Rearranging, we have

$$c_H \leq \ (1-\delta)c_L + \delta p .$$

Thus, for p near p_H, we have precisely the same condition as for adverse selection. The decision to choose high quality when it is a once-and-for-all decision is equivalent to the decision to choose high quality on an on-going basis.

There is a difference between the adverse selection and moral hazard cases. With adverse selection, the seller signals high quality by the introductory offer. Once signaled, the seller can sell at the monopoly price without loss of business since the product has been demonstrated as a high quality product. In contrast, with adverse selection, the seller is subject to ongoing incentive pressure, and a failure to sell a high quality good will result in a collapse of business.

6.3: Burning Money

Introductory offers have a flaw in a world in which both high quality and low quality goods are efficient for different types of customers. The low introductory price is intended to signal high quality, but instead may suggest that the product is actually the low quality product, which also now by hypothesis trades in equilibrium. Moreover, offering a high quality product at a low quality cost is likely to induce lots of buyers of the low quality product to switch to the subsidized high quality product, perhaps rendering the signaling unprofitable.

The introductory price offer depended on the following logic. The firm makes a price offer sufficiently low that a low quality firm would never imitate it. There is no need to spend this money on the consumer directly. Instead, the money can be spent on flashy advertisements, large buildings, or donations to public libraries. The point of signaling is to spend money so that it is impossible to recoup the expenditures unless the company gets repeat business. Any kind of conspicuous, irrevocable and irrecoverable expenditure will do.

A firm engaging in burning money would still charge a high price for the product. Doing avoids the risk of a bad signal about quality and deters customers who would be unwilling to pay the cost of the high quality good from consuming it.

7: The Coase Conjecture

In the paper "Durability and Monopoly," Nobel Laureate Ronald Coase proposes the startling hypothesis that the monopoly seller of a durable good will tend to price at marginal cost, absent some mechanism for committing to withhold supply. (Such mechanisms include leasing rather than selling, planned obsolescence, increasing marginal cost (which makes delay rational), and promises to repurchase at a fixed price.) The logic takes three steps. First, having sold the monopoly quantity at the monopoly price, the seller would like to sell a bit more, because the seller need not cut price on units already sold. Second, consumers will rationally anticipate such price cuts, and thus will hold out for future prices. Third, if the seller can change prices sufficiently fast, the path must go to marginal cost arbitrarily quickly, that is, the price will be marginal cost. This idea came to be known as the Coase conjecture.

Essentially the Coase conjecture holds that a monopolist compete with future incarnations of himself. Even though the most profitable course of action is to sell the monopoly quantity immediately, and then never sell again, the monopolist cannot resist selling more once the monopoly profit is earned. That is, subgame perfection condemns the monopolist to low profits.

7.1: The Commitment Solution

It is useful to consider the commitment solution as a benchmark, and to introduce notation. The seller's marginal cost is set to zero. Suppose time is discrete, with periods t=1,2,… Both the seller and the buyers discount each period at δ. Market demand is given by q, and is composed of a continuum of individuals.

The commitment solution involves a sequence of prices p_1, p_2,… This series of prices is non-decreasing without loss of generality, since no consumer will wait to buy at a higher price. A consumer with a value v will prefer time t to time t+1 if

(*) $v - p_t > \delta(v - p_{t+1})$

These equations define a sequence of critical values v_t that make the buyer indifferent between purchasing at t and purchasing at t+1. (Note that the incentive constraint on buyers shows that, if a buyer with value v chooses to buy before time t, then all buyers with values exceeding v have also purchased by this time.)

$$v_t - p_t = \delta(v_t - p_{t+1})$$

This set of equations can be solved for pt in terms of the critical values:

$$p_t = (1-\delta)v_t + \delta p_{t+1} = (1-\delta)v_t + \delta((1-\delta)v_{t+1} + \delta p_{t+2}) = \dots$$
$$= (1-\delta)\sum_{j=0}^{\infty} \delta^j v_{t+j}.$$

The monopolist sells $q(v_t)$-$q(v_{t-1})$ in period t, where v_0 is defined so that $q(v_0)$=0. The monopolist's profits are

$$\pi = \sum_{t=1}^{\infty} \delta^{t-1} p_t q_t = \sum_{t=1}^{\infty} \delta^{t-1}(1-\delta)\left(\sum_{j=0}^{\infty} \delta^j v_{t+j}\right)\left(q(v_t) - q(v_{t-1})\right)$$

$$= (1-\delta)\left[\sum_{t=1}^{\infty} q(v_t)\delta^{t-1}\left(\sum_{j=0}^{\infty} \delta^j v_{t+j}\right) - \sum_{t=1}^{\infty} q(v_{t-1})\delta^{t-1}\left(\sum_{j=0}^{\infty} \delta^j v_{t+j}\right)\right]$$

$$= (1-\delta)\left[\sum_{t=1}^{\infty} q(v_t)\delta^{t-1}\left(\sum_{j=0}^{\infty} \delta^j v_{t+j}\right) - \sum_{t=2}^{\infty} q(v_{t-1})\delta^{t-1}\left(\sum_{j=0}^{\infty} \delta^j v_{t+j}\right)\right]$$

$$= (1-\delta)\left[\sum_{t=1}^{\infty} q(v_t)\delta^{t-1}\left(\sum_{j=0}^{\infty} \delta^j v_{t+j}\right) - \sum_{t=1}^{\infty} q(v_t)\delta^{t}\left(\sum_{j=0}^{\infty} \delta^j v_{t+1+j}\right)\right]$$

$$= (1-\delta)\left[\sum_{t=1}^{\infty} q(v_t)\delta^{t-1}\left(\sum_{j=0}^{\infty} \delta^j v_{t+j} - \delta\sum_{j=1}^{\infty} \delta^{j-1} v_{t+j}\right)\right]$$

$$= (1-\delta)\sum_{t=1}^{\infty} q(v_t)\delta^{t-1} v_t = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} q(v_t) v_t .$$

Thus, the optimum level of v_t is constant at the one-shot profit maximizing level, which returns the profits associated with a static monopoly. The ability to dynamically discriminate does not increase the ability of the monopolist to extract rents from the buyers.

How does the requirement that the monopolist play a subgame perfect strategy affect the monopolist's profits? To simplify the analysis, let demand be linear: $q(p)=1-p$. Consider a game that ends at time T. Let a_t refer to the highest value customer remaining in the population at the end of time t, so that the set of values remaining at the beginning of time t is uniformly distributed on $[0,a_{t-1}]$, and the quantity purchased at time t is $a_{t-1} - a_t$.

In the last period, the monopolist is a one-shot monopolist, and thus charges the price $p_T=½\, a_{T-1}$ and earns profits $\pi_T = ¼ a_{T-1}^2$. This can be used as the basis of an induction to demonstrate that

$p_t = \lambda_t\, a_{t-1}$ and $\pi_t = \chi_t a_{t-1}^2$.

The last values satisfy $\lambda_T = ½$ and $\chi_T = ¼$.

The value a_t is determined by consumer indifference between buying at t and buying one period later, along with the beliefs that the monopolist will follow the equilibrium pricing pattern in the future, so that

$$a_t - p_t = \delta(a_t - p_{t+1}) = \delta(a_t - \lambda_{t+1} a_t),$$

or

$$p_t = a_t\,(1 - \delta + \delta\lambda_{t+1}).$$

Thus,

$$\pi_t = p_t\,(a_{t-1} - a_t) + \delta\pi_{t+1} = (1 - \delta + \delta\lambda_{t+1}) a_t\,(a_{t-1} - a_t) + \delta\chi_{t+1} a_t^2$$

Maximizing this expression over a_t, we see that the firm chooses p_t to induces a_t satisfying

$$a_t = \left(\frac{1 - \delta + \delta\lambda_{t+1}}{2(1 - \delta + \delta\lambda_{t+1} - \delta\chi_{t+1}} \right) a_{t-1}.$$

Feeding this expression into π_t and simplifying gives

$$\pi_t = \left(\frac{(1 - \delta + \delta\lambda_{t+1})^2}{4(1 - \delta + \delta\lambda_{t+1} - \delta\chi_{t+1}} \right) a_{t-1}^2.$$

We have, at this point, verified the induction hypothesis – p_t is linear in $a_{t\text{-}1}$ and π_t is quadratic, provided p_{t+1} is linear in a_t and π_{t+1} is quadratic

Since $\lambda_t\, a_{t\text{-}1} = p_t = a_t\,(1 - \delta + \delta\lambda_{t+1}) = \dfrac{(1 - \delta + \delta\lambda_{t+1})^2}{2(1 - \delta + \delta\lambda_{t+1} - \delta\chi_{t+1}} a_{t-1}$,

$$\lambda_t = \frac{(1 - \delta + \delta\lambda_{t+1})^2}{2(1 - \delta + \delta\lambda_{t+1} - \delta\chi_{t+1}},$$

and,

$$\chi_t = \frac{\pi_t}{a_{t-1}^2} = \frac{(1 - \delta + \delta\lambda_{t+1})^2}{4(1 - \delta + \delta\lambda_{t+1} - \delta\chi_{t+1}} = \frac{\lambda_t}{2}.$$

This permits the solution for λ_t in terms of λ_{t+1}.

$$\lambda_t = \frac{(1 - \delta + \delta\lambda_{t+1})^2}{2(1 - \delta + ½\delta\lambda_{t+1})}.$$

Define $f(\lambda) = \dfrac{(1-\delta+\delta\lambda)^2}{2(1-\delta+½\delta\lambda)}$, so that $\lambda_t = f(\lambda_{t+1})$. $f(0)=$ ½ $(1-\delta)$ and $f(1) = 1/(2-\delta)$. It is readily shown that f is increasing and strictly convex for $\delta \in (0,1)$. There is a unique fixed point for f, which occurs at

$$\lambda^* = \frac{\sqrt{1-\delta}-(1-\delta)}{\delta} < ½.$$

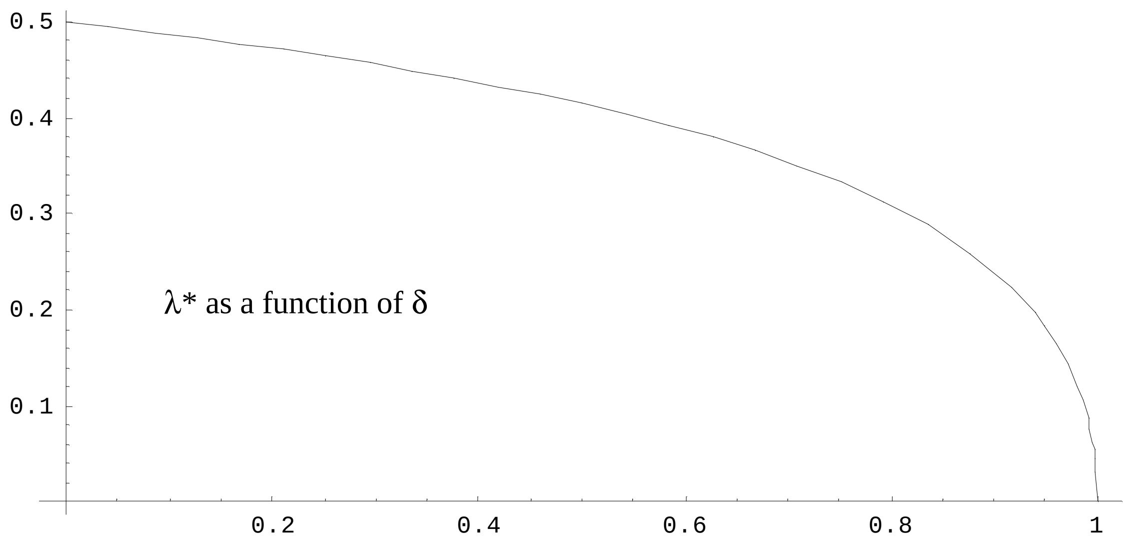

Since λ_T=½, the sequence λ_t is increasing in t to ½. For games with very large values of T, λ_1 is very close to λ^*. The opening price offered by the monopolist is λ_1, because a_0=1. The Coase conjecture amounts to the claim that, when the monopolist can cut prices very rapidly, the opening price is close to marginal cost, which was set to zero. The ability to cut prices very rapidly corresponds to a large discount factor – little discounting goes on between each pricing period. The Coase conjecture is in fact true, because

$$\lim_{\delta\to 1}\lambda^* = \lim_{\delta\to 1}\frac{\sqrt{1-\delta}-(1-\delta)}{\delta} \to 0.$$

This equilibrium is representative of all equilibria in the "gap" case, which is the case that arises when consumer valuation exceeds marginal cost by some positive amount for all consumers. In this case, price converges to the minimum consumer valuation, rather than to marginal cost. The "gap" case is not empirically relevant. The gap case corresponds to backward-induction equilibria because in fact the monopolist will sell all its output in finite time in equilibrium.

In addition, Gul, Sonnenschein and Wilson show that equilibria with stationary strategies also have the Coase property.

In contrast, Ausubel and Deneckere show that Coasian equilibria can be used to sustain other equilibria in which the monopolist makes positive profits. Consider a decreasing price path pt and suppose that consumers hold the conjecture that the monopolist will follow that path in equilibrium, and if that consumers see any price charged other than a price from the path, consumers believe that the monopolist will play the Coase path. The Coase path produces very low profits, so the threat of such beliefs is sufficient to sustain a large set of equilibria. These equilibria are a bit weird, since consumers beliefs help support prices that sustain high profits, but are sequential, if nonstationary, equilibria nevertheless. They are reasonable in that they predict a declining path of prices, with any lower prices suggesting even faster future decreases in prices.

7.2: Mitigating the Coase Problem

Do we really believe that a durable goods monopolist prices at marginal cost? There are several strategies a monopolist might employ to prevent his tendency to compete with himself.

1. Other equilibria don't have this property, but stationary (history independent) ones typically do.

2. Leasing vs. selling: A monopolist that leases the durable good no longer has the incentive to cut price to bring in new consumers, because he must cut the price to the old ones as well, who are, after all, leasing and can always turn it in and release at the lower price.

3. Return policy or money back guarantee: Suppose the monopolist allows one to return the good for the full purchase price, to be credited against future purchases. Then consumers need not wait for prices to fall - they've been given a guaranteed low price if prices do fall. This, of course, provides the same kind of disincentive to lower prices as leasing. Used by department stores on calculators in the early 1970s.

4. Destroy the production facility: used for limited edition items on occasion.

5. Make remaining in the market expensive: the profits in future periods are, of course, a decreasing function of time; the monopolist is cutting price and the high value consumers have left. Therefore, eventually, the monopolist will choose to exit [Singer almost exited the sewing machine industry on this logic, but stayed in after its announcement that it would leave because of public relations, or so they said]. However, this is mitigated by the entry (birth) of new consumers. But new consumers also mean that the monopolist has less incentive to cut prices, because he also cuts the price to the high value consumers. This can cause price cycles in some models, with the monopolist occasionally cutting price to grab the low value consumers who are accumulating.

6. Keep the marginal cost secret: If buyers aren't informed about the true marginal cost, their expectations will be influenced by the prices the monopolist charges, and the monopolist may not cut price, so as to convince consumers that he has high cost.

7. Planned Obsolescence: Note that a good that breaks down after one period - i.e. a nondurable good - is equivalent to leasing rather than selling. Thus, a monopolist who must sell has an incentive to reduce the durability of his good, that is, plan obsolescence. The standard example is a textbook manufacturer who frequently introduces new editions, to kill off the used book

market, and is therefore producing a nondurable good when a durable good is feasible and often optimal.

7.3: Capacity Choice

McAfee and Wiseman (2005) consider a monopolist who faces the Coase problem, but has a cost of producing quickly. To illustrate the problem of capacity choice, first consider a seller who sells in continuous time, but has a maximum production K per unit of time. Buyers are given by an inverse demand p. This seller will then sell Kt by time t. Let T be the time where the market is saturated, so that $p(KT)$=0; q_0 is the quantity KT, which may be infinite. Suppose the seller charges $P(t)$ at time t; a buyer with value $v = p(Kt)$ who buys at time s obtains profits of

$$u(Kt) = \max_s e^{-rs}\left(p(Kt) - P(Ks)\right).$$

By the envelope theorem:

$$u'(Kt) = e^{-rt} p'(Kt)$$

Thus, since $u(KT) = 0$,

$$e^{-rt}(p(Kt) - P(Kt)) = u(Kt) = u(KT) - \int_t^T u'(Ks)K\,ds = -\int_t^T e^{-rs} p'(Ks)K\,ds$$

$$= -\int_{Kt}^{KT} e^{-ry/K} p'(y)\,dy = e^{-rt} p(Kt) - \int_{Kt}^{KT} \frac{r}{K} e^{-ry/K} p(y)\,dy,$$

where the last equality follows from integration by parts and the observation that $p(KT) = 0$, and the second-to-last equality from a change of variable. Setting $z = Kt$, then, we obtain

$$P(z) = e^{rt} \int_z^{q_0} \frac{r}{K} e^{-ry/K} p(y)\,dy.$$

Thus, the firm's profits are

$$\pi(K) = \int_0^T e^{-rt} KP(Kt)\,dt = \int_0^{q_0}\int_z^{q_0} \frac{r}{K} e^{-ry/K} p(y)\,dy\,dz$$

$$= \int_0^{q_0} z \frac{r}{K} e^{-rz/K} p(z)\,dz \qquad \text{(integration by parts and substituting } p(q_0) = 0\text{)}$$

$$= \int_{0}^{q_0} \left(zp(z)\right)' e^{-rz/K} dz$$ (integration by parts and substituting $p(q_0) = 0$)

$$= \int_{0}^{q_0} MR(q) e^{-rq/K} dq$$ (definition of marginal revenue when $g = 0$).

This equation holds for any K, but we now consider what a monopolist who committed to K would choose. Note that the maximized value of profits is independent of the interest rate r, since a change in r is compensated by a change in K. Let $a = r/K$, and a^* maximize profits over K. Thus,

$$\pi = \max_a \int_0^{q_0} MR(q) e^{-aq} dq$$

$$= \max_a \int_0^{q_0} R(q) a e^{-aq} dq$$ (integration by parts, $R(q_0) = R(0) = 0$)

$$\geq \max_a \int_0^{q_m} R(q) a e^{-aq} dq$$ ($p(q)$ is decreasing)

$$= \max_a \int_0^{q_m} \frac{q}{q_m} R(q_m) a e^{-aq} dq\text{, so}$$

$$\frac{\pi}{R(q_m)} \geq \max_a \int_0^{q_m} \frac{q}{q_m} a e^{-aq} dq$$

$$= \max_a - e^{-aq_m} + \int_0^{q_m} \frac{1}{q_m} e^{-aq} dq$$ (integration by parts)

$$= \max_a \frac{1 - e^{-aq_m}(1 + aq_m)}{aq_m} = \gamma.$$

The value γ is approximately 0.298425. Thus, a monopolist with capacity commitment earns at least 29.8% of the profits earned by a perfectly price-discriminating monopolist. Moreover, even if capacity is very nearly free, McAfee and Wiseman show that the Coasian monopolist would never install capacity beyond the commitment level – the smallest positive cost of capacity insures that the monopolist earns at least 29.8% of the commitment capacity. The reason is that the monopolist correctly realizes that any additional capacity will influence the expectations of the buyers – the monopolist will be unable to restrain himself from using any capacity installed. Moreover, even if the buyers expect the monopolist to install additional capacity, he will not. The problem with expecting Coasian sales is that these require the monopolist to invest in capacity beyond the profit-maximizing level of capacity, at which point the monopolist is paying to harm profits. No profit-maximizer will engage in such behavior, hence the beliefs by the buyers that the monopolist is a profit-maximizing entity insure that the buyers don't expect unreasonable capacity increases; that is only possible when capacity is literally free.

However, costly capacity, once installed, is still used; it is an open question whether a monopolist will destroy capacity in order to hold up prices. The problem with the destruction of capacity is that it makes installation of too much capacity less costly, which might actually worsen the monopolist's profits.

8: Milgrom-Weber Auction Theory

Milgrom and Weber's 1982 paper on auctions and bidding remains one of the most impressive economics papers ever written. In order to generalize the solution to auction problems introduced by Vickrey, they developed new mathematics and statistics and then applied it to economics problems. This technology of complementarities has had dramatic applications in other economic environments. In this section, we examine their technology and see how it helps with the analysis of bidding schemes.

Auctions and bidding are important because they represent the foundation of market price formation. Generally we think of prices in markets as being established by some kind of bidding process; the theory developed here formulates rational bidding in environments with asymmetries of information about value and willingness to pay.

8.1: Complementarity

A building block of the theory is the notion of complementarity. Let $x \wedge y$, $x \vee y$, refer to the component-wise minima (*x meet y*) and maxima (*x join y*), respectively.

A function f:$\mathrm{R}^n \rightarrow \mathrm{R}$ is *supermodular* if $f(x \vee y) + f(x \wedge y) \geq f(x) + f(y)$.

Remark: If f is twice-differentiable, then supermodularity reduces to: $i \neq j \text{ implies } \frac{\partial^2 f}{\partial x_i \partial x_j} \geq 0$.

This, in turn, is equivalent to "increasing differences." That is, for $x_i > y_i$,

$f(x_1,..., x_{i-1}, x_i, x_{i+1},..., x_n) - f(x_1,..., x_{i-1}, y_i, x_{i+1},..., x_n)$ is non-decreasing in x_j for $j \neq i$.

If f is a payoff function and supermodular, the variables of f are said to be *complements*.

8.2: Affiliation

If the function log f is supermodular, f is said to be *log-supermodular*, sometimes abbreviated log-spm. If f is a density and is log-supermodular, then the random variables with density f are said to be *affiliated*. If there are only two such variables, f is said to have the *monotone likelihood ratio property* (MLRP).

The following six basic facts about affiliation and supermodularity are useful. I give simple arguments based on differentiability where possible, but the theorems hold more broadly.

(i) Affiliation is equivalent to the statement that $E[\alpha(\mathbf{X})|a_i \leq X_i \leq b_i]$ is non-decreasing in a_i, b_i for all non-decreasing functions α.

Proof: Consider $\varphi(y)$=E[$\alpha(x,Y)|Y=y, a \leq x \leq b$]. Below, expectations refer to conditioning on $Y=y$, $a \leq x \leq b$.

$$\varphi'(y) = E\left[\frac{\partial\alpha}{\partial y}\right] + \int \alpha(x,y)\left(\frac{f_y(s\mid y)}{\int_a^b f(z\mid y)dz} - \frac{f(x\mid y)\int_a^b f_y(z\mid y)dz}{\left(\int_a^b f(z\mid y)dz\right)^2}\right)dx$$

$$= E\left[\frac{\partial\alpha}{\partial y}\right] + E\left[\alpha(x,y)\frac{f_y(x\mid y)}{f(x\mid y)}\right] - E[\alpha(x,y)]\,E\left[\frac{f_y(x\mid y)}{f(x\mid y)}\right] = E\left[\frac{\partial\alpha}{\partial y}\right] + COV(\alpha, \frac{f_y}{f}).$$

If the MLRP is satisfied, this is nonnegative. Conversely, let α be increasing in x and constant in y. Then φ is non-decreasing for all y, a, and b if and only if the MLRP is satisfied.

(ii) Non-decreasing functions of affiliated r.v.'s are affiliated (see Milgrom-Weber for the proof; I know of no calculus-based proof).

Let x, y have density $f(x,y)$, and denote the density of y given x by $f_Y(y|x)$, with cdf $F_Y(y|x)$.

(iii) $F_Y(y|x)$ is non-increasing in x (First Order Stochastic Dominance).

The characteristic function of a set, 1_A, is the function which is 1 if $x\in A$ and 0 otherwise. Note $\Pr[X_i \geq x_i] = E[1_{\{X_i \geq x_i\}}]$. It follows that $\Pr[X_i \geq x_i \mid X_j = x_j]$ is non-decreasing in x_j.

(iv) f is log-supermodular if and only if f_Y is log-supermodular.

Proof: $\frac{\partial^2}{\partial x\,\partial y}\log\ f_Y(y\mid x) = \frac{\partial^2}{\partial x\,\partial y}\log\left(\frac{f(x,y)}{\int f(x,z)\,dz}\right)$

$$= \frac{\partial^2}{\partial x\,\partial y}\log\ f(x,y) - \frac{\partial^2}{\partial x\,\partial y}\log\left(\int f(x,z)\,dz\right) = \frac{\partial^2}{\partial x\,\partial y}\log\ f(x,y)$$

(v) Independently distributed random variables are affiliated. (Proof immediate.)

(vi) If $f(y|x)$ is log-supermodular, $F(y|x)$ is log-supermodular

Proof: $\frac{\partial}{\partial x}\frac{f(y\mid x)}{F(y\mid x)} = \frac{f_2(y\mid x)}{F(y\mid x)} - \frac{f(y\mid x)\,F_2(y\mid x)}{F(y\mid x)^2}$

$$= \frac{f(y\mid x)}{F(y\mid x)^2}\left[\frac{f_2(y\mid x)}{f(y\mid x)}F(y\mid x) - F_2(y\mid x)\right]$$

$$= \frac{f(y|x)}{F(y|x)^2}\left[\int_0^y \left(\frac{f_2(y|x)}{f(y|x)} - \frac{f_2(z|x)}{f(z|x)}\right) f(z|x)dz\right] \geq 0.$$

(vii) if f and g are log-supermodular, fg is log-supermodular. Proof is log fg = log f + log g.

8.3: Auction Environment

There is a single good for sale, and n bidders. Each bidder i privately receives a signal that is the realization of the r.v. X_i; the vector $(X_1,\ldots,X_n,\boldsymbol{S})$ are affiliated and the X_i's are symmetrically distributed. The payoff to bidder i is $u(X_i, X_{-i}, \boldsymbol{S})$. u is assumed non-decreasing in all arguments. All of the bidders know the environment; the only uncertainty for bidder i is X_{-i} and $\boldsymbol{S}$ and bidder i knows the joint distribution of the variables. We fix attention on bidder 1 and let $Y = \max\{X_2,\ldots,X_n\}$. Note that Y is affiliated with X_1. Let $f_Y(Y|x)$ be the density of Y given $X_1=x$, with distribution function F_Y. Let $v(x,y)=E[u|x_1=x,Y=y]$. Since u is non-decreasing, so is v.

8.4: Second Price Auction

In a second price auction, the high bidder obtains the object and pays the second highest bid.

A symmetric equilibrium bidding function is a function B_2 such that, given all other bidders bid according to B_2, the remaining bidder maximizes expected profit by bidding $B_2(x)$ given signal x. Consider bidder 1 with signal x who instead bids $B_2(z)$. This bidder earns

$$\pi = \int_0^z (v(x,y) - B_2(y))\, f_Y(y|x)dy.$$

In order for B_2 to be an equilibrium, π must be maximized at $z=x$, which implies

$$B_2(x) = v(x,x).$$

It is straightforward to show that B_2 is indeed an equilibrium, and is the only symmetric equilibrium.

If a reserve price (minimum acceptable bid) r is imposed, bidders with signals below x_r, where $E[v(x_r,Y)|Y\leq x_r]=r$, do not submit bids; otherwise the equilibrium is unperturbed. Note however, that the minimum submitted bid, $B_2(x_r) > r$!

Suppose the seller knows S_i. Should the seller tell the bidders S_i? The thought experiment is a commitment to honestly reveal the information prior to discovering it. Let

$$w(x,y,s) = E[u \mid X_1=x, Y=y, S_i=s].$$

$$v(y,y) = E[w(X_1,Y,S_i) \mid X_1 = Y = y]$$

$$= E[w(Y,Y,S_i) \mid X_1 = Y = y]$$

$$\leq E[w(Y,Y,S_i) \mid X_1 \geq Y = y].$$

The seller's revenue with no disclosure, R_N, is

$$R_N = E[v(Y,Y) \mid X_1 \geq Y]$$

$$\leq E[E[w(Y,Y,S_i) \mid X_1 \geq Y] \mid x_1 > Y]$$

$$= E[w(Y,Y,S_i) \mid X_1 > Y] = R_I,$$ the revenue with disclosure of S_i.

That is, it always pays to reveal accurate information, rather than conceal it.

The thought experiment of revealing accurate information is sensible if the seller can commit to actually revealing accurate information. That is not always possible, and when the seller can't commit not to lie about the information, bidders will likely ignore the information since if bidders condition on it, the seller will misstate it. In some settings, where the seller can be penalized for lying (fraud) or when the seller repeatedly sells something, the seller has the option of committing to telling the truth, and the theorem says that this is better than silence.

8.5: First Price Auction

In a first price auction, the high bidder obtains the object and pays her bid. Suppose B_1 is a symmetric equilibrium. The profits to bidder 1, with signal x, who bids $B_1(z)$, are:

$$\pi = \int_0^z (v(x,y) - B(z)) f_Y(y \mid x) dy.$$

Maximizing with respect to z, and setting z=x, yields the first order differential equation

$$B_1'(x) = \frac{f_Y(x \mid x)}{F_Y(x \mid x)} (v(x,x) - B_1(x)).$$

Suppose that the reserve price is zero. Then the differential equation has solution

$$B_1(x) = \int_0^x e^{-\int_y^x \frac{f_Y(z|z)}{F_Y(z|z)} dz} \frac{f_Y(y \mid y)}{F_Y(y \mid y)} v(y,y) dy.$$

(If the reserve price $r > 0$, the screening level is x_r and B_1 satisfies $B_1(x_r)$=r.) Integrating $B_1(x)$ by parts and assuming v(0,0)=0, we have:

$$B_1(x) = v(x,x) - \int_0^x e^{-\int_y^x \frac{f_Y(z|z)}{F_Y(z|z)}dz} \left(\frac{d}{dy} v(y,y)\right) dy.$$

Conditional on winning with a signal of x (probability $F_Y(x|x)$), a bidder in a second price auction pays

$$EB_2 = \int_0^x v(y,y) \frac{f_Y(y\,|\,x)}{F_Y(x\,|\,x)} dy = v(x,x) - \int_0^x \frac{F_Y(y\,|\,x)}{F_Y(x\,|\,x)} \left[\frac{d}{dy} v(y,y)\right] dy.$$

Note that $\log F_Y(x\,|\,x) - \log F_Y(y\,|\,x) = \int_y^x \frac{f_Y(z\,|\,x)}{F_Y(z\,|\,x)} dz \geq \int_y^x \frac{f_Y(z\,|\,z)}{F_Y(z\,|\,z)} dz$ (by(vi))

and thus, $\frac{F_Y(y\,|\,x)}{F_Y(x\,|\,x)} \leq e^{-\int_y^x \frac{f_Y(z|z)}{F_Y(z|z)}dz}$.

Since v is non-decreasing, the expected payment by a winning bidder with signal x is higher in a second price auction than in a first price auction.

9: Bibliography

Adams, W.J. and J. L. Yellen, “Commodity Bundling and the Burden of Monopoly,” *Quarterly Journal of Economics*, 1976, 475-98.

Ausubel, Larry and Raymond Deneckere, “Reputation in Bargaining and Durable Goods Monopoly,” *Econometrica* 57, 1989, 511-32.

Boiteaux, Marcel, “La Tarification des Demandes en point: application de la theorie de la vent au cout marginal,” *Revue Generale de l’Electricite* 58, August, 1949, 321-340.
(English translation available in J *Business* 33, April 1960.)

Butters, G., "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies* 44, 465-91 (1980).

Carlson, John and R. Preston McAfee, “Joint Search for Several Goods”, *Journal of Economic Theory 32*, April 1984.

Coase, Ronald, "Durability and Monopoly," *Journal of Law and Economics* 15: 143-149, 1972.

Dana, Jr., James D. 1999a, ‘Equilibrium price dispersion under demand uncertainty: the roles of costly capacity and market structure’, *RAND Journal of Economics*, vol. 30, no. 4 (Winter), pp. 632-660.

Deneckere, Ray and R. Preston McAfee, “Damaged Goods,” *Journal of Economics and Management Strategy 5*, no. 2, Summer, 1996, 149-74.

Diamond, Peter, A Model of Price Adjustment, Journal of Economic Theory,1971, 156-68.

Gallego, Guillermo, and van Ryzin, Garrett 1994, ‘Optimal dynamic pricing of inventories with stochastic demand over finite horizons’, *Management Science*, vol. 40, no. 8 (August), pp. 999-1020.

Gul, Faruk, Hugo Sonnenschein & Robert Wilson, “Foundations of Dynamic Monopoly and the Coase Conjecture,” *Journal of Economic Theory* 39, 1986, 155-190.

Hong, Pilky, Ashish Nayyar and R. Preston McAfee, “Equilibrium Price Dispersion with Consumer Inventories,” *Journal of Economic Theory*, 105 n.2, August 2002, 503-17.

Kahn, Charles, “The Durable Goods Monopoly and Consistency with Increasing Costs”, *Econometrica* 54, March 1986, 275-294.

McAfee, R. Preston and Vera te Velde, “Dynamic Pricing in the Airline Industry”, *Handbook on Economics and Information Systems*, Ed: T.J. Hendershott, Elsevier, forthcoming 2006.

McAfee, R. Preston, John McMillan and Michael Whinston, “Commodity Bundling by a Monopolist,” *Quarterly Journal of Economics*, May 1989, 371-83.

McAfee, R.P. “Endogenous Availability, Cartels and Merger in an Equilibrium Price Dispersion,” *Journal of Economic Theory* 62, no. 1, February 1994, 24-47.

McAfee, R. Preston and Thomas Wiseman, “Capacity Choice Counters the Coase Conjecture,” unpublished.

McAfee, R. Preston and John McMillan, “Auctions and Bidding,” *Journal of Economic Literature*, Summer, 1987.

McAfee, R. Preston and Daniel Vincent, “The Price Decline Anomaly,” *Journal of Economic Theory* 60, June, 1993, 191-212.

McAfee, R. Preston and John McMillan, “Bidding Rings,” *American Economic Review* 82, no.3, June 1992, 579-99.

McAfee, R. Preston, “Pricing Damaged Goods,” unpublished manuscript, http://vita.mcafee.cc, 2006.

McAfee, R. Preston, “The Price is Right Mysterious,” *Engineering and Science*, 3, 2005.

McAfee, R. Preston, *Competitive Solutions: The Strategist’s Toolkit*, Princeton, NJ: Princeton University Press, 2002.

McAfee, R. Preston, “Coarse Matching,” *Econometrica*, vol. 70, 2002, issue 5, 2025-2034.

Milgrom, Paul and Robert Weber, “A General Theory of Auctions and Bidding,” *Econometrica*, November 1982.

Milgrom, Paul, 1989, “Auctions and Bidding: A Primer,” *Journal of Economic Perspectives* 3: 3-22.

Mussa, Michael and Rosen, Sherwin 1978, ‘Monopoly and product quality’, *Journal of Economic Theory*, vol. 18, pp. 301-317.

Oren, Shmuel, Stephen Smith and Robert Wilson, “Capacity Pricing,” *Econometrica* 53, May 1985, 545-566.

Pesendorfer, Wolfgang, “Fashion Cycles,” *American Economic Review*, 1995.

Stigler, George, “The Economics of Information,” *Journal of Political Economy* 69, June 1961, 213-25.

Stokey, Nancy L. 1979, 'Intertemporal price discrimination', *The Quarterly Journal of Economics*, vol. 93, no. 3 (August), pp. 355-371.

Talluri, Kalyan T. and van Ryzin, Garrett J. 2004, *The Theory and Practice of Revenue Management,* Kluwer Academic Publishers, Dordrecht, Netherlands.

Tirole, Jean, *Industrial Organization*, 1985, 66-72, 153-158.

Varian, Hal, "A Model of Sales," *American Economic Review*, 70, 651-659 (1980).

Varian, Hal, "Price Discrimination and Social Welfare." *American Economic Review*, 75(4):870-875, 1985.

Williamson, Oliver, "Peak Load Pricing and Optimal Capacity under Indivisibility Constraints," *American Economic Review*, 56, September 1966, 810-827.

Wilson, Robert B., *Nonlinear Pricing*,

Wilson, Robert, "Efficient and Competitive Rationing," *Econometrica* 57, January 1989.

10: Appendix 1: Derivations

Proof that $q_n(t) = \frac{(\beta(0)-\beta(t))^{K-n} B_n(t)}{(K-n)!B_K(0)}$.

Suppose the theorem is true at $n \le K$. Then we must verify that

$$q'_{n-1}(t) = \frac{B'_{n-1}(t)}{B_{n-1}(t)} q_{n-1}(t) - \frac{B'_n(t)}{B_n(t)} q_n(t)$$

Plugging in the hypothesized values:

$$q_{n-1}(t) = \frac{(\beta(0)-\beta(t))^{K-n+1} B_{n-1}(t)}{(K-n+1)!B_K(0)}, \text{ so}$$

$$q'_{n-1}(t) = -\beta'(t)\frac{(\beta(0)-\beta(t))^{K-n} B_{n-1}(t)}{(K-n)!B_K(0)} + \frac{(\beta(0)-\beta(t))^{K-n+1} B'_{n-1}(t)}{(K-n+1)!B_K(0)}$$

But

$$q'_{n-1}(t) = \frac{B'_{n-1}(t)}{B_{n-1}(t)} q_{n-1}(t) - \frac{B'_n(t)}{B_n(t)} q_n(t)$$

$$= \frac{B'_{n-1}(t)}{B_{n-1}(t)} \frac{(\beta(0)-\beta(t))^{K-n+1} B_{n-1}(t)}{(K-n+1)!B_K(0)} - \frac{B'_n(t)}{B_n(t)} \frac{(\beta(0)-\beta(t))^{K-n} B_n(t)}{(K-n)!B_K(0)}$$

$$= B'_{n-1}(t)\frac{(\beta(0)-\beta(t))^{K-n+1}}{(K-n+1)!B_K(0)} - B'_n(t)\frac{(\beta(0)-\beta(t))^{K-n}}{(K-n)!B_K(0)}$$

Thus, we have verified the formula if

$$-\beta'(t)\frac{(\beta(0)-\beta(t))^{K-n} B_{n-1}(t)}{(K-n)!B_K(0)} + \frac{(\beta(0)-\beta(t))^{K-n+1} B'_{n-1}(t)}{(K-n+1)!B_K(0)}$$
$$= B'_{n-1}(t)\frac{(\beta(0)-\beta(t))^{K-n+1}}{(K-n+1)!B_K(0)} - B'_n(t)\frac{(\beta(0)-\beta(t))^{K-n}}{(K-n)!B_K(0)}$$

or

$$\beta'(t)\frac{(\beta(0)-\beta(t))^{K-n}B_{n-1}(t)}{(K-n)!B_K(0)}=B_n'(t)\frac{(\beta(0)-\beta(t))^{K-n}}{(K-n)!B_K(0)}$$

or

$$\beta'(t)B_{n-1}(t)=B_n'(t)\,.$$