Versioning

Carl Shapiro
Hal R. Varian*
University of California, Berkeley

September 1998 (revised: September 19, 1998)

1 A simple numerical example

We've described several dimensions on which you might version your product. The goal over versioning is to get your customers to sort themselves into groups with different values for your product; you can adjust both the design of the product and its price in order to influence this sorting.

A simple numerical example will illustrate just how this works. In this example, "delay" is the dimension along which you have decided to version, but the same principles work for all the dimensions described above. Suppose that you have 100 customers, of which 40 are "impatient" and 60 are "patient." The impatient customers will pay $100 for immediate information, but only $40 for delayed information; timeliness is worth an extra $60 to the these customers. In contrast, the patient customers will pay $50 for immediate information and $30 for delayed information; timeliness is only worth an extra $20 to these customers. Table 1 summarizes the customers and their values.

*Notes to accompany *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, 1998. Material on second-degree price discrimination adopted from Hal R. Varian, *Microeconomic Analysis*, 3rd edition, 1992, W. W. Norton & Co.

	Impatient customers' value	Patient customers' value
Immediate version	100	50
Delayed version	40	30
Number of customers	40	60

Table 1: Willingness-to-pay for immediate or delayed information, along with number of each type.

Let's look at the revenues you can generate using various strategies for pricing your information in this setting. If you offer just the immediate version, your best approach is to set a price of $50, sell to all 100 customers, and earn total revenues of $5000. This is clearly better than setting a price of $100, selling only to the 40 impatient customers, and earning revenues of $4000. Offering only a delayed version is clearly worse than this: why degrade the product for everyone, and force your own price down as a result? It follows that the best you can do with a single-version, single-price strategy is to earn $5000.

Of course, you would love to sell the immediate version to everyone, charging the impatient customers $100 and the patient customers $50, if you could manage to do it. This would give you revenue of $40 \times 100 + 60 \times 50 = 7000$. This is the most you could possibly hope to earn, as it involve maximizing total value and then extracting *all* of the value from each type. Selling the immediate version to everyone, but at different prices, is precisely the perfect price discrimination we discussed in the previous chapter. However, to implement this sort of strategy, you need a way of determining whether any particular customer is impatient (and thus should be charged the full $100) or is patient (and thus suitable for the $50 rate). If the impatient customers all work at investment houses and the patient customers teach in schools, your problem is solved.

But what happens if you can't identify the high-value and low-value users based on observable characteristics such as type of business, age, location, or gender? Then you need to resort to versioning. What prices should you charge for the two immediate and delayed versions?

The first thing to try would be to target each version at one group, for a price equal to that group's willingness to pay. In particular, you could try to sell the immediate version for \$100 and the delayed version for \$30. Faced with these offering, the patient consumers will indeed buy the delayed version. But the impatient consumers will not buy the immediate version, as you have planned. If an impatient consumer buys the immediate information at \$100 his value is \$100 and the cost is \$100, so his net benefit is zero. If he buys delayed information, its value to him is \$40 but he only has to pay \$30, so he gets a net benefit of \$10. This means that the delayed version is a better deal for the impatient consumer. As a result, pricing the immediate version at \$100 and the delayed version at \$30 will not induce customers to self-select as planned: the impatient consumers will not select the immediate version, which was meant for them, but will instead choose the cheaper delayed version.

So what do you do? The trick is to *discount* the immediate information sufficiently so that the impatient consumers will indeed buy it. The net benefit they receive from purchasing the immediate version has to be the same (or just slightly more) than the surplus they could get by purchasing the delayed version, which is \$10. This means you can't charge more than \$90 for the immediate version, if you charge \$30 for the delayed version. At these prices the 60 patient consumers will buy the delayed version, generating revenues of \$1800, and the 40 impatient consumers will buy the immediate version, generating revenues of \$3600. Your total revenues from versioning will be \$5400. This is not as good as pricing based on identity, which yielded \$7000, but it's still a lot better than the \$5000 you would get from selling only one version.

2 Second-degree price discrimination

This is a mathematical analysis based on the treatment in Hal R. Varian, *Microeconomic Analysis*, 3rd edition, 1992, W. W. Norton & Co. We initially examine the case of pricing based on quantity, then, at the end, show how this can be interpreted as a quality discrimination model.

Suppose that there are two potential consumers with utility functions $u_i(x) + y$, for $i = 1, 2$. For simplicity, normalize utility so that $u_i(0) = 0$. Consumer i's maximum willingness-to-pay for some consumption level x will be denoted by $r_i(x)$. It is the solution to the equation

$$u_i(0) + y = u_i(x) - r_i(x) + y.$$

The left-hand side of the equation gives the utility from zero consumption of the good, and the right-hand side gives the utility from consuming x units and paying a price $r_i(x)$. By virtue of our normalization, $r_i(x) \equiv u_i(x)$.

Another useful function associated with the utility function is the marginal willingness-to-pay function, i.e., the (inverse) demand function. This function measures what the per-unit price would have to be to induce the consumer to demand x units of the consumption good. If the consumer faces a per-unit price p and chooses the optimal level of consumption, he or she must solve the utility maximization problem

$$\max_{x,y} \quad u_i(x) + y \tag{1}$$

$$\text{such that} \quad px + y = m. \tag{2}$$

As we have seen several times, the first-order condition for this problem is

$$p = u_i'(x).$$

Hence, the inverse demand function is given explicitly by 2: the price necessary to induce consumer i to choose consumption level x is $p = p_i(x) = u_i'(x)$.

We will suppose that the maximum willingness-to-pay for the good by consumer 2 always exceeds the maximum willingness-to-pay by consumer 1; i.e., that

$$u_2(x) > u_1(x) \text{ for all } x.$$

We will also generally suppose that the *marginal* willingness-to-pay for the good by consumer 2 exceeds the marginal willingness-to-pay by consumer 1; i.e., that

$$u_2'(x) > u_1'(x) \text{ for all } x.$$

Thus it is natural to refer to consumer 2 as the **high demand** consumer and consumer 1 as the **low demand** consumer.

We refer to consumer 2 as the *high-demand* consumer and consumer 1 as the *low-demand* consumer. The assumption that the consumer with the larger total willingness-to-pay also has the larger marginal willingness-to-pay is sometimes known as the single crossing property since it implies that any two indifference curves for the agents can intersect at most once.

3 Analysis

Suppose that the monopolist chooses some (nonlinear) function $p(x)$ that indicates how much it will charge if x units are demanded. Suppose that consumer i demands x_i units and spends $r_i = p(x_i)x_i$ dollars. From the viewpoint of both the consumer and the monopolist all that is relevant is that the consumer spends r_i dollars and receives x_i units of output. Hence, the choice of the function $p(x)$ reduces to the choice of (r_i, x_i). Consumer 1 will choose (r_1, x_1) and consumer 2 will choose (r_2, x_2).

The constraints facing the monopolist are as follows. First, each consumer must want to consume the amount x_i and be willing to pay the price r_i:

$$u_1(x_1) - r_1 \geq 0 \quad (3)$$
$$u_2(x_2) - r_2 \geq 0. \quad (4)$$

This simply says that each consumer must do at least as well consuming the x-good as not consuming it. Second, each consumer must prefer his consumption to the consumption of the other consumer.

$$u_1(x_1) - r_1 \geq u_1(x_2) - r_2 \quad (5)$$
$$u_2(x_2) - r_2 \geq u_2(x_1) - r_1. \quad (6)$$

These are the so-called self-selection constraints. If the plan (x_1, x_2) is to be feasible in the sense that it will be voluntarily chosen by the consumers, then each consumer must prefer consuming the bundle intended for him as compared to consuming the other person's bundle.

Rearrange the inequalities in the above paragraph as

$$r_1 \leq u_1(x_1) \quad (7)$$
$$r_1 \leq u_1(x_1) - u_1(x_2) + r_2 \quad (8)$$
$$r_2 \leq u_2(x_2) \quad (9)$$
$$r_2 \leq u_2(x_2) - u_2(x_1) + r_1. \quad (10)$$

Of course, the monopolist wants to choose r_1 and r_2 to be as large as possible. It follows that in general one of the first two inequalities will be binding and one of the second two inequalities will be binding. It turns out that the assumptions that $u_2(x) > u_1(x)$ and $u_2'(x) > u_1'(x)$ are sufficient to determine which constraints will bind, as we now demonstrate.

To begin with, suppose that (9) is binding. Then (10) implies that

$$r_2 \leq r_2 - u_2(x_1) + r_1,$$

or

$$u_2(x_1) \leq r_1.$$

Using (2) we can write

$$u_1(x_1) < u_2(x_1) \leq r_1,$$

which contradicts (7). It follows that (9) is not binding and that (10) is binding, a fact which we note for future use:

$$r_2 = u_2(x_2) - u_2(x_1) + r_1.$$

Now consider (7) and (8). If (8) were binding, we would have

$$r_1 = u_1(x_1) - u_1(x_2) + r_2.$$

Substitute from (3) to find

$$r_1 = u_1(x_1) - u_1(x_2) + u_2(x_2) - u_2(x_1) + r_1,$$

which implies

$$u_2(x_2) - u_2(x_1) = u_1(x_2) - u_1(x_1).$$

We can rewrite this expression as

$$\int_{x_1}^{x_2} u_1'(t)\, dt = \int_{x_1}^{x_2} u_2'(t)\, dt.$$

However, this violates the assumption that $u_2'(x) > u_1'(x)$. It follows that (8) is not binding and that (7) is binding, so

$$r_1 = u_1(x_1).$$

Equations (3) and (3) imply that the low-demand consumer will be charged his maximum willingness-to-pay, and the high-demand consumer will be charged the highest price that will just induce him to consume x_2 rather than x_1.

The profit function of the monopolist is

$$\pi = [r_1 - cx_1] + [r_2 - cx_2],$$

which upon substitution for r_1 and r_2 becomes

$$\pi = [u_1(x_1) - cx_1] + [u_2(x_2) - u_2(x_1) + u_1(x_1) - cx_2].$$

This expression is to be maximized with respect to x_1 and x_2. Differentiating, we have

$$\begin{aligned} u_1'(x_1) - c + u_1'(x_1) - u_2'(x_1) &= 0 \qquad (11) \\ u_2'(x_2) - c &= 0. \qquad (12) \end{aligned}$$

Equation (11) can be rearranged to give

$$u_1'(x_1) = c + [u_2'(x_1) - u_1'(x_1)] > c,$$

which implies that the low-demand consumer has a (marginal) value for the good that exceeds marginal cost. Hence he consumes an inefficiently small amount of the good. Equation (12) says that at the optimal nonlinear prices, the high-demand consumer has a marginal willingness-to-pay which is equal to marginal cost. Thus he consumes the socially correct amount.

Note that if the single-crossing property were not satisfied, then the bracketed term in (3) would be negative and the low-demand consumer would consume a *larger* amount than he would at the efficient point. This can happen, but it is admittedly rather peculiar.

The result that the consumer with the highest demand pays marginal cost is very general. If the consumer with the highest demand pays a price in excess of marginal cost, the monopolist could lower the price charged to the largest consumer by a small amount, inducing him to buy more. Since price still exceeds marginal cost, the monopolist would make a profit on these sales. Furthermore, such a policy wouldn't affect the monopolist's profits from any other consumers, since they are all optimized at lower values of consumption.

4 A graphical treatment

The price discrimination problem with self-selection can also be treated graphically. Consider Figure 1 which depicts the demand curves of the two consumers; for simplicity we assume zero marginal cost. Figure 1 A depicts the price discrimination if there is no self-selection problem. The firm would simply sell x_h^o to the high-demand consumer and x_l^o to the low-demand consumer at prices that are equal to their respective consumer's surpluses—i.e., the areas under their respective demand curves. Thus the high-demand consumer pays $A + B + C$ to consume x_h^o and the low-demand consumer pays A to consume x_l^o.

However, this policy violates the self-selection constraint. The high-demand consumer prefers the low-demand consumer's bundle, since by choosing it he receives a net surplus equal to the area B. In order to satisfy the self-selection constraint, the monopolist must offer x_h^o at a price equal to $A + C$, which leaves the high-demand consumer a surplus equal to B no matter which bundle he chooses.

This policy is feasible, but is it optimal? The answer is no: by offering the low-demand consumer a slightly smaller bundle, the monopolist loses the profits indicated by the black triangle in Figure 1 B, and gains the profits indicated by the shaded trapezoid. Reducing the amount offered to the low-demand consumer has no first-order effect on profits since the marginal willingness-to-pay equals zero at x_l^o. However, it increases profits non-marginally since the high-demand consumer's willingness-to-pay is *larger* than zero at this point.

At the profit-maximizing level of consumption for the low-demand consumer, x_l^m in Figure 1 C, the marginal *decrease* in profits collected from the low-demand consumer from a further reduction, p_1, just equals the marginal *increase* in profits collected from the high-demand consumer, $p_2 - p_1$. (Note that this also follows from equation (3).) The final solution has the low-demand consumer consuming at x_l^m and paying A, thereby receiving zero surplus from his purchase. The high-demand consumer consumes at x_h^o, the socially correct amount, and pays $A + C + D$ for this bundle, leaving him with positive surplus in the amount B.

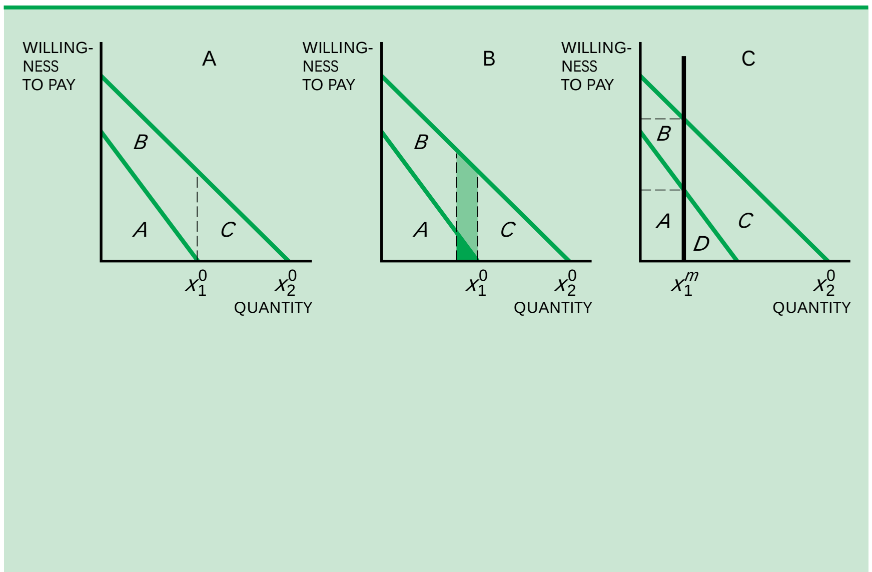

Figure 1: Second-degree price discrimination. Panel A depicts the solution if self-selection is not a problem. Panel B shows that reducing the bundle of the low-demand consumer will increase profits, and panel C shows the profit-maximizing level of output for the low-demand consumer.