The Economics of Software Products: an Example of Market Failure

Carlos, del Cacho

Elca Informática

29 March 2009

The Economics of Software Products: an Example of Market Failure

Abstract

In this paper we examine pricing imperfections in software companies by analyzing the case of Microsoft, and we uncover the presence of pervasive dead-weight losses derived from the inability of the producer to achieve first degree price discrimination. Because the nature of software is such that it can be reproduced an infinite number of times at practically zero cost once the first copy is manufactured, the amount of these losses in terms of efficiency can be substantial, which opens the door for external intervention in the market. We finish by suggesting a simple policy rule in this direction, although the applicability may be limited to the theoretical realm, as it can distort the incentives of private enterprise as a provider of software products.

****

As has been described by Varian, the production methods of certain types of goods are characterized by large fixed costs and comparatively low variable costs, conferring them important economies of scale. In these cases, pricing at marginal cost may not recoup all the production costs, consequently firms should sometimes engage in differential pricing to ensure a socially desirable outcome, and regulations forcing flat pricing may end up hurting economic efficiency as a result.

Some industries, like the telecommunications sector, lend themselves to this task better than others do. An unfortunate example of this latter kind is a large segment of the software industry, where arbitrage by third parties cannot be easily prevented, and the ability of the producer to price discriminate is therefore severely diminished. This applies to medium to low priced products with high volumes of sales using the Internet or retail shops as distribution channels. It is not applicable however to low volume high priced items, typically supported by sales force personnel, where segmentation is indeed possible and copyright protection can be adequately enforced. For the rest of this document we will dissect the former type of information goods.

Software as a product is indeed a very interesting economic problem to focus on, because the marginal costs of production, or distribution for that matter, are rather insignificant, so in a perfect world the supply side should be able to meet nearly all potential demand. By demand in this context I mean any customer willing to buy at a price that satisfies both his wants and those of the producer, i.e. making a healthy profit on each and every sale.

The peculiarities of software also indicate that maybe it should be a publicly produced good, in the sense that it is perhaps better managed by the government than by private means. I will try to make a case for placing Microsoft Office in public hands, and find out that there is potential for an acquisition where everybody involved wins.

Let's imagine for a second that we are in *Economic Wonderland*, where all sorts of miracles happen, rather than in the real world. In this imaginary place, the producer has perfect information about the demand curve, and is able to identify how much each and every customer is willing to pay for the product it tries to sell, setting the price individually for each and every customer willing to buy. It is assumed that our customer

will always buy if the *perceived value* received by purchasing the product exceeds the price the producer sets, or if the *perceived value* is equal to said price, it will be indifferent about the purchasing choice altogether and just as likely to acquire the product or not.

However, just a quick glance reveals that we don't live in that imaginary world. Many wrong economic policies in the past were based on oversimplifications like this one, stemming from assumptions coming way back from neoclassical economics. And well, no software company in the world can achieve perfect market segmentation, unless they tried to auction the product, which is basically impossible for software because it lacks the natural scarcity of other auctionable goods, like for example, art pieces or search engine advertising space. What's more, because the decision makers of our software company have really no idea about the willingness to buy of each individual customer, they try to make educated guesses through market research, forecasting based on past sales data, etc.

And perhaps much more importantly, even in the unlikely event that they did manage to gather all that information, they could still not achieve perfect market segmentation. Why so? Because customers in the real world interact and communicate with each other, and they would sooner or later find out that they're being deceived when realizing that they paid a certain amount and someone else paid much less for what's essentially the same commodity. Potential customers, aware of the whole process, may then defer purchases expecting to find a better deal somewhere else, or even decide not to buy at all. The brand name of Microsoft would likely be permanently affected in the marketplace.

I briefly searched for a reliable source of information on profit margins of Microsoft Office but couldn't find any; the data I found suggests it might be around 70-80%, which is sufficiently high for the sake of this argument. As to why it is this high, the explanation can be found in vendor lock-in through proprietary formats, and the network effects present in the natural exchanges of data in said formats between different owners of the product. Once the product achieves critical mass, vendor lock-in basically forbids the entrance of new competitors to the network of existing buyers, and the network effects present in document exchanges consolidate a monopolistic market position.

Ignoring the recurring support costs and marketing costs of a product based company, there is an opportunity to buy it out based on the inability of Microsoft to meet all market demand that they really would want to reach to maximize their profit, because as we saw before, they cannot discriminate between customers willing to pay 150$ or customers willing to pay 200$ for the same product unless they use very complicated pricing mechanisms, which won't work in practice. Let's say, Microsoft decides to set its price for Microsoft Office at 200$ because their analysis shows it maximizes their revenue. Based on their 70-80% profit margin, they would also be willing to sell it at 150$, 100$, or 75$, if they only could.

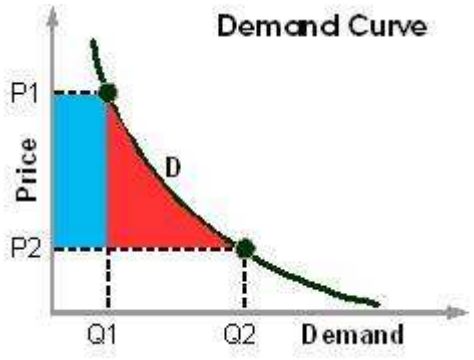

Figure 1: Consumer surplus gain when pricing at P2

That means that, we could theoretically make a case to buy out the Microsoft Office product line and place it in public hands based on the following observation: it is actually worth more to the public than to Bill Gates himself, hence there is an opportunity for a buy out where all parties involved win. Microsoft's inability to achieve perfect price segmentation is no problem at all for the government, which can simply charge a token fee to cover the recurring support costs, since the profit margin of Microsoft was 70-80%, the government must charge at least 20-30%, let's say price P2 in figure 1, of the amount that Microsoft charged initially (P1), to cover the operating costs. Now, this figure is just a ballpark estimate that may or may not be accurate because the increased demand satisfied by the government charging lower prices will certainly increase support costs, but perhaps not in a linear fashion, and because software scales rather nicely with volume, it may be actually lower in the end.

Furthermore, we also have to consider that buying out the market leader may also push other producers out of the market, and this may have other positive economic benefits in the sense that production of inferior substitute goods is actively discouraged. Given the nearly zero marginal costs of producing and distributing a software product, this is really a good thing, because the demand for Microsoft Office will increase and the costs will drop as a result of these rather peculiar properties, namely that software scales with volume. The *consumer surplus* from going from price P1 to P2 is increased, meaning that the difference between what the customer pays and the received *perceived value* is higher.

By how much, you ask? By the area marked in figure 1. The area in red is extra benefit accrued to the customers now that was not reachable in the previous scenario, because Microsoft was unable to capture it either. Conversely, the area marked in blue in the diagram is value that has switched hands from the producer to the consumer. This is value that Microsoft was capturing before but it has been displaced now with the change in pricing.

So how exactly did this happen? We saw before that Microsoft is very unlikely to ever achieve perfect pricing on its own, so if Microsoft calls the price, it will leave out of the market customers that would have otherwise bought. As it stands, it is far less costly to

broadcast the price to all potential customers than to get into your prospects' brains to know how much they are really willing to pay. I would even suggest that the second option is impossible, but I cannot foresee all the technological advancements that science may offer us within some years.

An economist would say that the transaction costs for both options are asymmetrical. It is precisely this asymmetry that allows the customers to self-select themselves and reach a more efficient outcome in the second scenario than in the first one. In theory, it is possible that all present and future interested buyers, right up to the end of the product life cycle, those that can be reached by Microsoft right now as well as those that have and will be left out of the market because of imperfections in pricing, could get together and buy out the whole Microsoft Office product line and share it among themselves, as they can relinquish part of their extra *consumer surplus* and hand it to the producer.

On the one hand, the producer will gladly accept the arrangement as this increases the amount of profit that it would otherwise be able to capture. And on the other hand, the consumers as a group win as long as they retain part of their newly attained *consumer surplus* during the negotiation.

As interesting as it may be, this scenario won't ever happen in practice, because while the consumers as a group may win, they are likely to try to misrepresent their real intentions of buying to each other, resulting into what is known as the *free rider problem*. They have incentives to state that they are only willing to pay up to an amount when they would be really willing to go higher. This occurs because consumers compete with each other just as much as producers do to acquire goods or services in the economy, so they may be very tempted to take such a course of action when given the chance.

Let us analyze now this point more closely. In figure 2 we see the matrix of expected payments for all possible strategies that can be followed when agreeing on the purchase. We say that a potential client collaborates with the rest if he reveals his real intentions of buying to the group, making a deposit into a common fund for the lowest amount between P1 and the maximum sum of money he is willing to spend when acquiring the product. If every member were to be completely honest, they could compute the revenue obtained by Microsoft in the original situation, i.e. the sum of the contributions of the individuals with willingness to pay of at least P1. In this scenario they would make an offer to Microsoft slightly above this value and the extra consumer surplus that they obtain they share it among themselves (the remaining sum of money in the fund after the purchase), let us say, in equal slices (X).

Group \ Individual	Collaborates	Does not collaborate
Rest collaborates	X	X + Y
Rest does not collaborate	0	0

Figure 2: Matrix of payments for the individual strategies when faced with the purchase

It is to be expected that some potential clients will lie about their willingness to pay. If they are only but a few and the rest of the group collaborates, they would be able to retain for themselves an extra value Y resulting from the difference between the maximum price they would be willing to pay and the one they revealed to the group. As we said that they were just a few, the distributed consumer surplus (X) decreases though not much, and the difference we can consider as negligible.

However, if many of them act dishonestly, the estimate derived when adding up the contributions from the members with willingness to pay of at least P1 will be incorrect, and the deal may finally fall through, as Microsoft will reject it because the offered quantity is insufficient to their eyes. Each individual has an incentive to lie and thus obtain more benefit than the rest from the arrangement, but if enough individuals act that way, nobody gains anything. If we observe the matrix of payments we can see that not collaborating is a dominant strategy for every individual, because the expected value of said strategy is at least as good as collaborating, with the potential of being better if the rest of the group collaborates but the individual in question does not. We say in this case that we face a Nash equilibrium and the agreement will not be reached in the end.

So perhaps the government should buy out the Microsoft Office product line, and price on cost to maximize welfare. In any case, carrying out a detailed cost-benefit analysis for such a project would be excruciating, as collecting revenue through taxation imposes burdens of its own in the economy and their effects are difficult to quantify. In addition, the government would be subsidizing future Microsoft Office customers at the expense of other citizens in its country, so this alternative is to be discouraged. The people that benefit from buying out Microsoft Office should foot the bill for themselves, but as we already said, they probably can't because they are in direct competition for other goods and services in the economy.

Hold on a second, didn't we just discover a rather serious market inefficiency? This is a theoretical win-win proposition for everybody involved that would lead to what an economist would call a Pareto efficient outcome, but it's just not taking place. I guess this can only indicate that markets are at least a bit more inefficient than some may think.

Economic efficiency is maximized if any of the goals below is achieved…

- The producer is able to perfectly segment the market and price accordingly (the *producer surplus* is maximized).

- The price covers just the costs and the profit is zero (the *consumer surplus* is maximized).

- The producer manages to always price below the maximum price a consumer is willing to stand (then neither of them is maximized). We will focus on the first two cases and I'll leave it as an exercise to the reader to find out why we shouldn't care at all about this third case.

Competition between consumers helps to achieve the first goal (*producer surplus maximization*), but as we saw before, it is practically impossible to reach this outcome in the real world, so we can immediately rule it out as a viable path to efficiency,

because when the producer calls the price, due to the high transaction costs to reach the prospects, it will necessarily leave out some potential customers that may have otherwise bought.

On the contrary, competition between producers helps to achieve the second goal (*consumer surplus maximization*), and this has far more interesting implications. First of all, no producer is willing to take up an investment just to break even. Hence, reaching this second outcome perfectly within a competitive framework is also impossible. And secondly, a modest amount of profit may not be worthwhile if the capital needed can be put to work elsewhere for a higher return on investment.

But what this whole thing really means is that high profit margins should really be actively suppressed whenever possible by regulators, ensuring that competition between producers brings the *producer surplus* down to reasonable levels, and at the same time increasing the *consumer surplus* and helping the formerly unmet demand to enter the market.

A profit maximizing firm is unlikely to pursue this on its own as it conflicts with its inherent interests, so Adam Smith's invisible hand may perhaps need a little bit of help from time to time to reach a more efficient outcome.

Let's go back to Microsoft Office... Antimonopoly laws may be passed on whenever a market exhibits characteristics conductive to a permanent market advantage resulting in an ever lasting monopoly. For Microsoft Office, these are vendor lock-in and network effects as mentioned before. A possible path of action to improve economic efficiency is to force Microsoft to open all their proprietary formats so that other competitors can enter the market and thus erode profit margins.

As far as the new competitors improve efficiency through lower prices more than they raise total cost, competition is a desirable, albeit an imperfect way to remedy the situation. Typically the existing group of sellers will lower prices until they face inelastic demand whereby their revenue cannot be increased by further lowering their price, until the high profit margins alert other entrepreneurs and they enter the specific business niche, thus eroding those same margins in their quest for market share.

And given that software as a product is self-induced scarcity in the land of plenty, in an even more radical attempt to improve efficiency, we could envision the government taking full control of Microsoft Office by effectively expropriating the software at production costs, and to avoid the subsidy problem, taking out a loan to finance the purchase. Microsoft's shareholders would obviously be made worse off since they wouldn't receive the dividends coming from the profit stream, but if the government decided to price at a level just enough to repay the loan and cover ongoing support costs, everybody else would be better off.

Some may object that private property is a universal right and oppose such a measure from an ethical standpoint, but I might as well argue that Bill Gates is rich enough already and he shouldn't mind that much. From an economic point of view, this policy would result in a more efficient result than the present state of being, except for the fact that software piracy already achieves more or less the desired effect without the need for intervention.

What do we mean by that? Many economic policies are shoved down the throats of innocent souls based on what is known as the Kaldor-Hicks efficiency criterion. This basically means that, given two possible resource allocation choices leading to two different outcomes, we can evaluate which is better by looking at the global picture rather than at each and every individual affected. Some people may win, some people may lose, others may be indifferent about the whole thing, but as long as the winners gain more than the losers lose, they might somehow manage to compensate them for their loss.

Well, here's the news, attacking the private property rights of Microsoft in this manner is also Kaldor-Hicks efficient. Thus pure *laissez faire* economics is wrong, or so we may be inclined to think. Of course, Microsoft Office would never come into being in the first place if the incentives to its production were to be removed completely, which leads us to ponder if software should be publicly produced instead, or at least part of it. As a consequence, the selective impulse and adoption of open source initiatives by the administration is not too farfetched, as this reduces the extent of the efficiency losses coming from the private sector.

Nonetheless, since I hear Hayek screaming from his grave, we need not be this revolutionary after all. Microsoft enjoys a monopolistic position although at the same time it faces fierce competition, not from the outside, but from its very own headquarters in Seattle. The nature of its business model is such that it receives a lump sum for each copy of Microsoft Office it sells, but in order to generate additional revenue, it is being forced to either reach new customers all the time, or given that software does not wear out like other goods, to improve its product offering so as to compel existing customers to buy a fresh license again of its refurbished releases. The existence of this second constraint pushes Microsoft to remain competitive even in the absence of external pressure.

However, it tries to escape from this loop by striking distribution agreements with hardware providers, bundling its new operating systems and software when a consumer acquires a notebook or a PC, as hardware does wear out and gets obsolete in a few years, inducing its customer base to upgrade to the new Microsoft software even if they hadn’t considered the possibility on their own. This practice should perhaps demand further attention from antitrust regulators as it tends to hinder competition in the software industry, although given that one firm can supply the whole market and new entrants won’t increase total output, competition might not be a desirable trait if it raises total industry costs faster than it reduces the losses in efficiency coming from imperfections in pricing.

In the world of physical goods, the entrance of competitors translates directly into a global increase of supply, but for the case of software products this premise does not hold true, because the supply of each individual vendor is potentially infinite, as software is a non-rival good. What’s more, every vendor replicates a core set of features available across all products in order to reach a marketable state, and this can be regarded as a waste of productive resources. Copying the functionality of a product of the complexity of Microsoft Office can easily range in the order of millions of dollars. OpenOffice, an open source alternative, has entered the scene by reverse engineering Microsoft’s file formats, but in order to do so it has also paid dearly in development costs.

Just so you get a glimpse of how bad the situation really is, let's leave Microsoft Office and its insurmountable barriers to entry aside, to focus on markets that don't offer that shield of protection against competition. Take for example the antivirus, backup, and CD replication utilities. These are horizontal solutions, therefore the size of their respective markets is huge, and there are no network effects here that prevent the entrance of new vendors. As a result, the fields are fairly competitive.

A quick search in one of the most popular software download sites in the Internet yields the following results:

Keyword	Results
antivirus	272 hits
cd burner	59 hits
backup	149 hits

Figure 3: Products found per key phrase

Each of these products solves a group of tasks in certain ways, some of these solutions are common among all the remaining products, and some of them are unique to each product and come as a result of innovation on the part of their creators, giving them a relative advantage over the products that lack them. If there is any proven market value in these differentiating features, the rest of the pack rapidly reacts by copying them and adding them to their respective products. In doing so, they also incur in most of the R&D costs of the innovator, and after a while the features no longer bring a competitive advantage.

Some countries, although not all, protect innovation by granting patents on software artifacts, but these are often difficult to enforce in court and they can easily be circumvented through minor variations if the scope of the patent is not ample enough.

Shall an antivirus firm find a better way to serve the needs of its customers, the 271 remaining competitors will eventually follow suit and replicate the winning combination in their own products. Rather than having it developed once, it will be developed again and again. If instead of competing between them, all the workers from all antivirus firms decided to collaborate to create only one antivirus product, the surplus workers could be deviated towards other productive occupations in the benefit of the economy as a whole. Unlike in other fields of human endeavour, competition breeds inefficiency in the software industry. However, as we will explain later, a Pareto improvement is possible.

Let us have a closer look. First we will consider vendor A alone in the marketplace, which faces the demand curve for the whole market. Because most of the cost is concentrated in the first few units, partially through fixed costs and partially through variable costs incurred when producing the first unit, the average cost curve can be regarded as a monotonically decreasing function, like the one shown in figure 4. This has a few interesting implications, that we proceed to unravel.

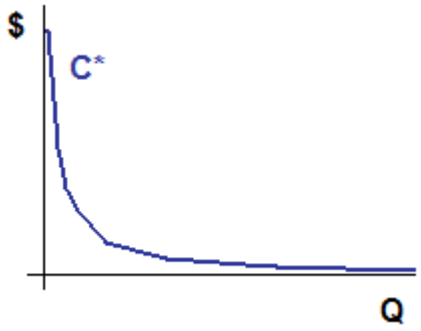

Figure 4: Average cost of software as a function of produced units

Now we take into account the effect of introducing competitor B with an identical product to that of our initial vendor. Figure 5 shows a joint demand curve, which is the quantity demanded at each pricing point provided that both products have the same price and move in unison. For every price, the joint demand is exactly half of that of the market demand, because both producers share the total revenue equally. This is a useful mental construct because whichever equilibrium point the market reaches, it must necessarily fall within this line, as states where both products have different prices cannot be regarded as stable.

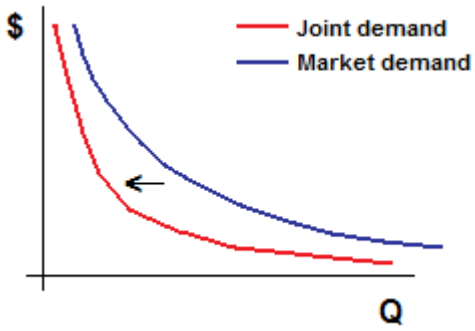

Figure 5: Joint demand curve vs. market demand curve

If the product from vendor A is priced below product from vendor B it will serve the whole market, and vendor B has an incentive to lower its pricing point and set it to the same price as product A, or lower, where A will be tempted to repeat the process or risk losing all customers. However, profit of both firms is maximized if they set their price at the level of unit demand elasticity, and although strictly speaking this is not a Nash equilibrium as both firms have an incentive to lower their price to gain the whole market, fear of retaliation is likely to leave the price unaltered and a tacit agreement is reached whereby the price remains above what a competitive outcome might suggest, i.e. where the profit is zero.

Notice that when both vendors share a piece of the market, the pricing point that makes the profit equal zero is higher than when one vendor is alone in the market. This happens because of the decreasing nature of the average cost function, and it means that given a market of limited size, software is a natural monopoly, that is, a firm can obtain above normal profits during a sustained period of time. A firm alone in the market is guaranteed to obtain profits if it sets its price in the range shown in figure 6, delimited by p1 and p2.

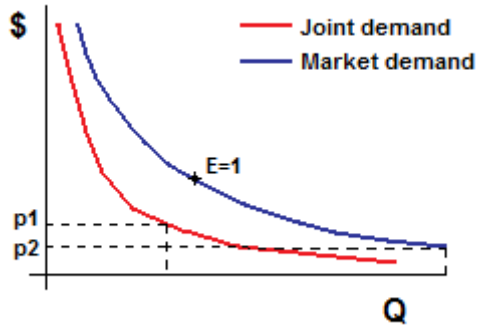

Figure 6: The market of software as a natural monopoly

We can also conceive that because a product with features A+B+C, where arbitrage cannot be avoided, has only one pricing point, there is a market opportunity for other firm to introduce a product that has just features B+C or C+D or any other combination whatsoever that serves a business purpose. That said, it will never pay off for a profit maximizing firm to introduce a product with functionality A+B+C if the original vendor sets the price within the price point specified in the range as it will result in losses for the company.

Nonetheless, we have observed before that software markets without substantial barriers to entry have plenty of competition. Since a product typically has a limited lifespan, a firm may decide to set its price above the range we mentioned if it obtains enough profits overall to offset the future losses of profit attained by attracting new entrants later on.

This fact, coupled with imperfect information on the side of the consumer, whereby he may be aware of only a limited set of products when deciding to purchase one of them, explains that there are 272 antivirus products in the market instead of just one. This has rather tragic consequences from an efficiency consideration, but we can theoretically improve on this outcome. Total industry cost increases linearly with each new entrant, which opens up the following possibility: an external benevolent regulator could theoretically buy ALL companies in the marketplace above the market price, i.e. considering their present and future profit stream, and arrive at a nearly Pareto improvement in overall efficiency.

It is quite logical that there is nothing we can do about sunk costs that the industry has incurred in producing the available array of products in the market, but we can improve

the outcome if we reduce future costs that are about to take place. An observation must be made: for the particular case of software, any market equilibrium is not a Nash equilibrium in itself, as a firm has an incentive to add additional functionality to its existing product and earn above normal profits. As we said before, the competitors will copy said functionality if it has market value, which they will immediately feel in the form of a loss in market share if they remain oblivious to it without doing anything about it.

Suppose we are a governmental entity and we can estimate the profit earned by each company in the antivirus marketplace during the whole product lifecycle, for example, ten years. If we offer 5% over that amount to the shareholders for their respective products and they act rationally, meaning that they attempt to maximize their profit, they will accept the offer. How do we recover this money? First of all, we can continue selling the products at the current market price and earn all but the 5% we overpaid for. But we can improve on this outcome if we consider the savings in development costs.

For the particular case of software, R&D cost pales in comparison with the cost of implementing and testing to deliver a real product. We could continue developing the leading product in the market and discontinue development for everything else altogether. Imagine development costs eat up 60% of the revenue for a given product. If we fire all the developers except for those of one company and give them 2 years salary as severance package for doing nothing (that they can use to retrain themselves for a different occupation where they will actually do something productive), we would be saving 47.8% of total industry costs, that can pay for the additional 5% in profit and then some more. We cannot ascertain for sure that this is a Pareto improvement because it is not clear if the developers will find a better job or not, but we give them a fair chance of doing so, enough to consider that perhaps they will.

Conclusion

In this paper we presented an overview of the economic principles that apply in the software industry. While it is widely accepted since a long time ago that a non discriminating monopoly results in non-efficient outcomes, the effect of competition in the software industry from the point of view of efficiency has been largely overlooked. Software products always result in inefficiencies either through the market power of a few players or through the wasted productivity of many small parties.

In the light of these findings, it would be tempting to conclude that the government must necessarily step in to improve the situation, but an indiscriminate attack on the private sector would likely stifle innovation and slow economic progress. Caution must then be exercised in extracting any normative implication out of this knowledge.

References:

Differential Pricing and Efficiency, Varian H. 1996

http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/473/394