

An introduction to Revenue sharing

Kris Van Bruwaene
VRT

Sharing the revenue between content providers and network operators is a recurring problem in media business modelling. It is generally a matter of negotiation between the parties, but those negotiations often lack a clear basis and have limits. This introduction will try to determine the limits of “fair, reasonable and non-discriminatory” revenue sharing, based on a simplified model. It does not presuppose any economic expertise from the reader.

The media value chain

Should a broadcast network provider pay the content provider for exploiting that content, or should the content provider pay the network operator for bringing that content to the customers? That is the problem we are dealing with, and the main cause of it is the media value chain (*Fig. 1 bottom*). In a traditional commercial value chain (*Fig. 1 top*) products are provided by manufacturers to retailers, who in turn distribute them to the end users. Money travels in the opposite direction, each party taking a reward for the value it adds to the product. The “just” reward (i.e. not overly exceeding the costs) is ensured by competition at every level of the value chain.

A typical characteristic of the media value chain is the existence of parallel flows of money that bypass the distributors. This is obviously the case with income from advertising, but also with licence fee and government aid. Even in many instances of pay television, the viewer has a direct commercial relation with the pay-TV operator, not with the network operator. This situation causes confused negotiations within the value chain, and often results in distorted competition, e.g. by exclusive deals or vertical integration between content providers and networks. This value chain is not unique to broadcasting, but also applies to other media, print in particular. Like in commerce and other media, there is a high mutual dependency between the value chain compo-

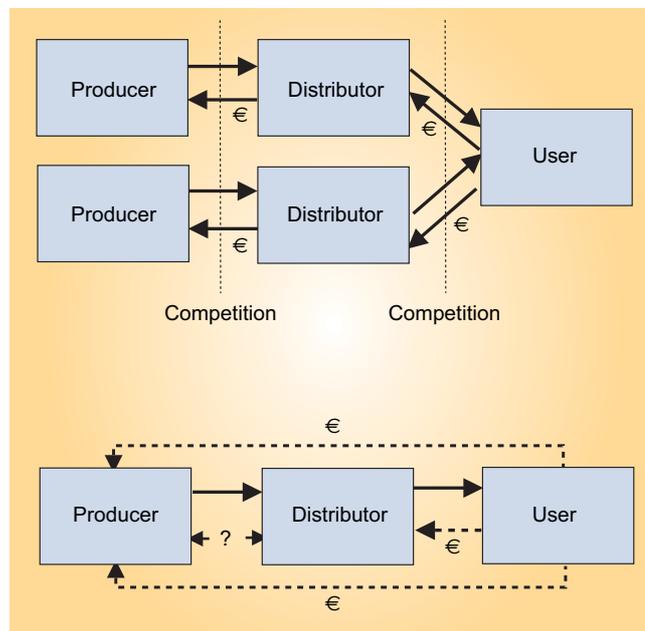


Figure 1
Top: traditional value chain, Bottom: media value chain. Cross links between various producers and distributors have been omitted for clarity

nents: networks without content will not attract customers, content without means of distribution is equally worthless.

Fig. 2 presents the media value chain in a more useful way. The right side of the horizontal axis shows what the customers pay directly to the network operator. The parallel money flows that depend on the market share are shown at the left side. These are typically advertising revenue, commissions on sales through home shopping and betting channels, and pay-television revenue arriving directly at the content provider; other parallel revenues from licence fee or income from other (e.g. international) markets do not depend on market share, although they can be important. They are considered fixed and not taken into account, as are the fixed costs. It is convenient to express revenues as an amount per customer and per time period (e.g. month, quarter or year).

Market size or penetration is represented on the vertical axis. The demand curve is well approximated by a downward sloping straight line. We will deal with the real demand curve later on.

We consider network operators and content providers in a *closed market* and with only *fixed costs*, by which we mean independent of the number of customers. This is not too unrealistic: the network incurs no extra costs when someone switches on the television set, nor does the production of a programme become cheaper when less people view it. In other words: there are high fixed costs, but the *marginal costs* per customer are zero. As a consequence, profit is directly related to revenue. The best point of operation for the network provider is therefore halfway along the demand line, where the product of the number of customers and the amount paid, represented by the revenue rectangle, reaches its peak.

The area to the left of the vertical axis represents the parallel revenue, arriving at the content provider without passing through the network operator. The upper triangle underneath the demand line represents the *unserved audience*, people that find the offer too expensive. The lower triangle is usually called the *consumer surplus*: customers in this triangle would be willing to pay more for the content and have the impression that they receive more than they pay for. *Welfare* is defined as the sum of total direct and parallel revenue of content and network providers and the consumer surplus; in other words, everything except the revenue lost by the unserved audience. Strictly speaking there is no relation between this economic definition of welfare and the meaning of the term in everyday use. It is however not forbidden to link the revenue part of welfare with employment, or attribute a social value to distributing content for free.

Free to air

Consider the case where content providers and network operators are united in a single company, that tries to maximize its revenue (Fig. 2), the so called *vertical integration*. When a media market has a high advertising potential and a lower willingness to pay, the best operating point can reach the top of the demand line, as shown in Fig. 3.

The content can then be offered for free, it is entirely paid for by the parallel revenue. This is a very desirable situation, in that it provides the

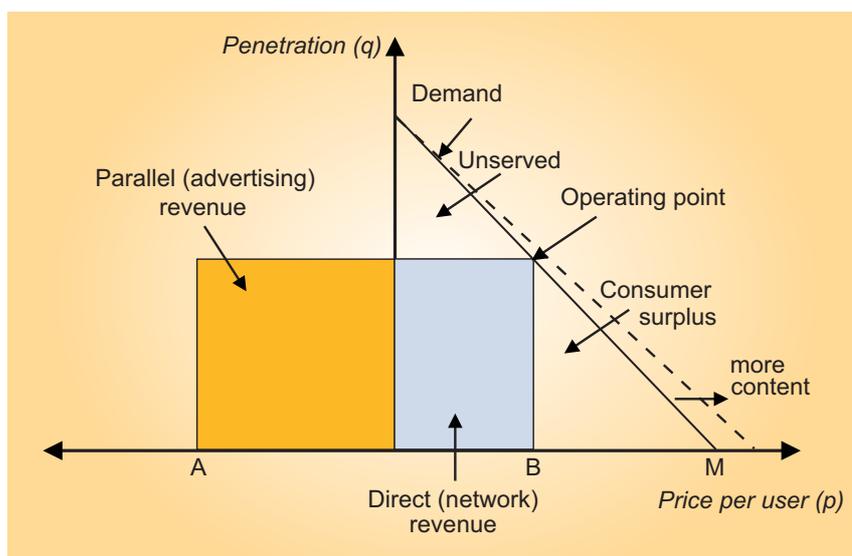


Figure 2

Revenue model for content providers and network operators in a closed market with fixed costs only. A represents the advertising spent per media user, M the maximum willingness of a user to pay. For a vertically-integrated company, the optimum price point lies halfway between A and M: the network operator reaches maximum revenue at $B = M/2$ and the content provider at $B = 0$ (free to air). The latter is also the point with the highest welfare.

highest possible consumer surplus and welfare. It has been the traditional model in broadcasting for decades, and still is in many countries, in both analogue and digital, terrestrially and on satellite, and in particular for radio. Moreover it is an economically stable point, where the price equals the marginal cost, both being zero. Several real world models exist, from the case where every content provider operates its own network, to a single network operator being paid by all content providers.

The main drawback of the free-to-air model is its limited content offer, due to funding: neither advertising nor licence fee can grow indefinitely. Spectrum scarcity is sometimes quoted as a limiting factor, but for television this can only have been the case until the arrival of broadcasting satellites. For radio however, spectrum may limit the capacity in densely populated areas.

In this model the network is paid by the content providers. The minimum payment is the network operating cost, which can be relatively low as there is neither customer administration (billing, smart card distribution etc.) nor marketing involved. The maximum fair payment is the highest revenue that the network operator could possibly pull from the market in a pay model, i.e. the right rectangle in *Fig. 2*. As can be noticed in *Fig. 3*, this amounts to a quarter of the parallel revenue. But excess revenue or growth in revenue had better be invested in more or better content rather than in the network, as this increases the willingness to pay – and consequently – welfare.

This successful model is currently under threat. The main breach was the advent of pay television in the 1980s. It was thought that a segment of the population was willing to pay for extra channels with exclusive content, and this turned out to be the case, although perhaps not as successful as many hoped for. We will delve deeper into this subject later on. Pay television as such does not undermine the free-to-air model, but by introducing multichannel networks and by increasing the willingness to pay for television, it laid the foundations for its potential demise.

The utility model

Let us now consider the typical European cable companies, as they exist in the Benelux countries, Germany, Switzerland and Scandinavia. Many of these originated as community antenna systems, offering the viewers a single access to a collection of various free-to-air programmes at a low cost. Many eventually evolved into larger scale utility companies. In the beginning, the content providers were neither paid nor paying for this (re)distribution of their content. This arrangement is beneficial to both viewers and content providers: it enlarges or at least retains the content provider's reach, thus maintaining or raising its commercial income, while at the same time increasing the choice for viewers at a cost not much higher than that of an individual aerial. Total welfare is somewhat reduced compared to the free-to-air model, but this may well be compensated by the larger number of programmes offered to the end users; that factor has not been considered in the present analysis.

After some consolidation phases, the reach and penetration of these cable companies become interesting for new commercial content providers – both free-to-air and pay-tv. Being able to feed their

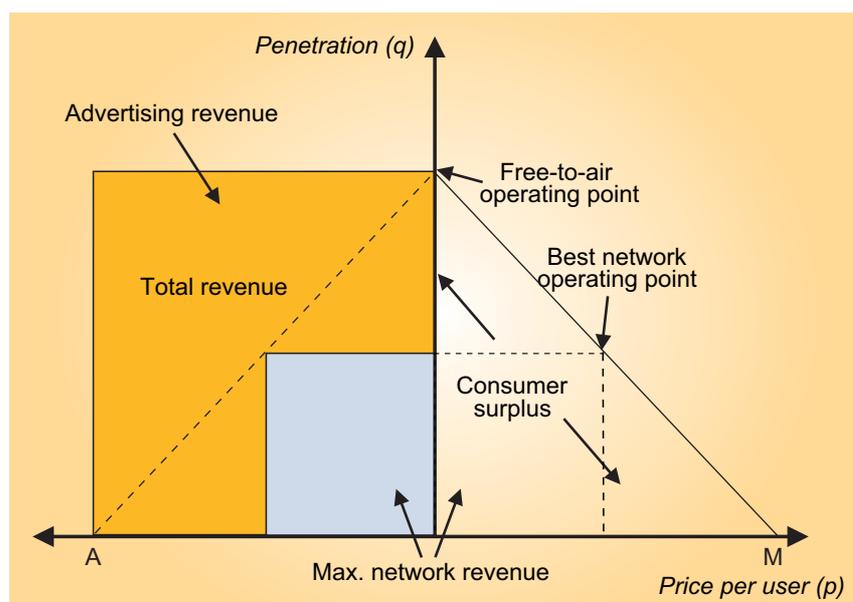


Figure 3
The free to air model is viable when the mid-point between A and M lies at the origin (as shown) or to the left of it

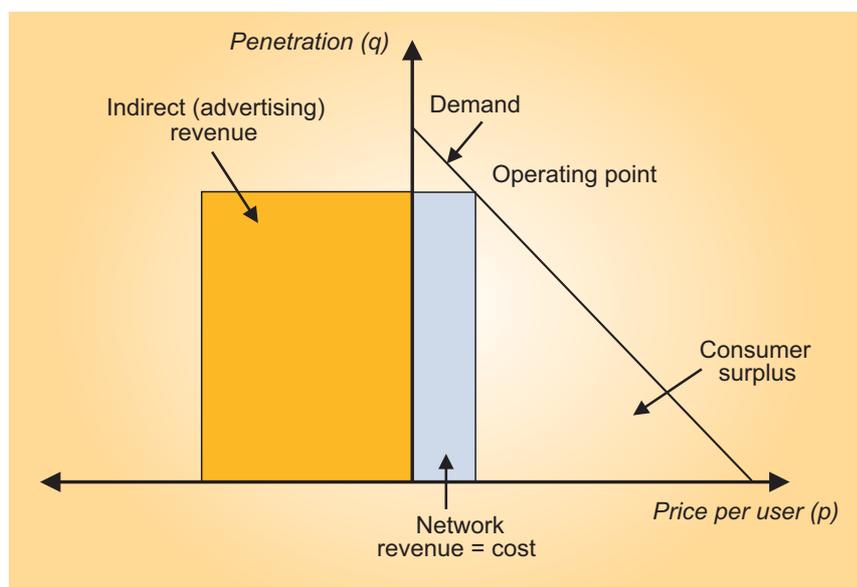


Figure 4
Utility model

signal directly into the cable networks is a clear bonus for these content providers, so one might expect them to pay the cable companies for their transportation. As long as there is enough capacity in the cable network, the extra money earned by the cable company can be used to lower the subscription fee for the end users or to expand the network, thus increasing reach and market penetration. The presence of these exclusive cable television channels further enhances the attractiveness of cable.

This utility model is illustrated in *Fig. 4*. The main advantage for

content providers is that their distribution is now paid entirely or mainly by the end users, so the money saved on distribution can be invested in content. It is a mild form of pay television. Cable TV is a local monopoly, and like most monopolies it tends to migrate its operating point towards the more comfortable middle of the demand curve. This would enlarge the unserved audience and is neither desirable for the users nor for the content providers that depend on advertising. Therefore the authorities should exercise a price cap, which is not uncommon for monopoly utility companies.

Revenue sharing is a simple matter in the utility model, in fact there need not be any sharing: the network costs can be paid entirely by the end users, the content production is covered by the parallel revenue of the content providers. Still, revenue sharing does occur frequently, in either direction. Usually content providers pay a fixed transportation fee to the network provider and in turn receive rights payments, depending on their market share. The net result often is a net income for content providers with a significant market share, and a net transfer to the network for small content providers. There is nothing wrong with this arrangement, it reflects demand and supply.

The vertically integrated optimum (*Fig. 2*) is reached when the additional income from a price increase at the network side exactly compensates the loss of advertising income at the content side. Although it has no relation with the utility operating point, based on the network costs, both points can lie close to each other in a market with limited advertising income.

Multiple networks and exclusivity

What happens when several networks compete in the same market? The demand curve for each of them is lowered to their market share as shown in *Fig. 5*. Supposing all operators have their own network, their operating point has to shift to the right in order to cover their network costs. This means that their service will become more expensive than in the utility-monopoly situation: in this case competition does not work in favour of the consumer. Even in the free-to-air model, the content providers will have to contribute towards the costs of all networks, leaving less money for content production.

Similarly, when a content provider strikes an exclusive deal with one network, it reduces its reach and lowers its potential revenue from advertising. It will try to compensate this by demanding a higher share of the network revenue, leading towards a pay television model.

Pay television

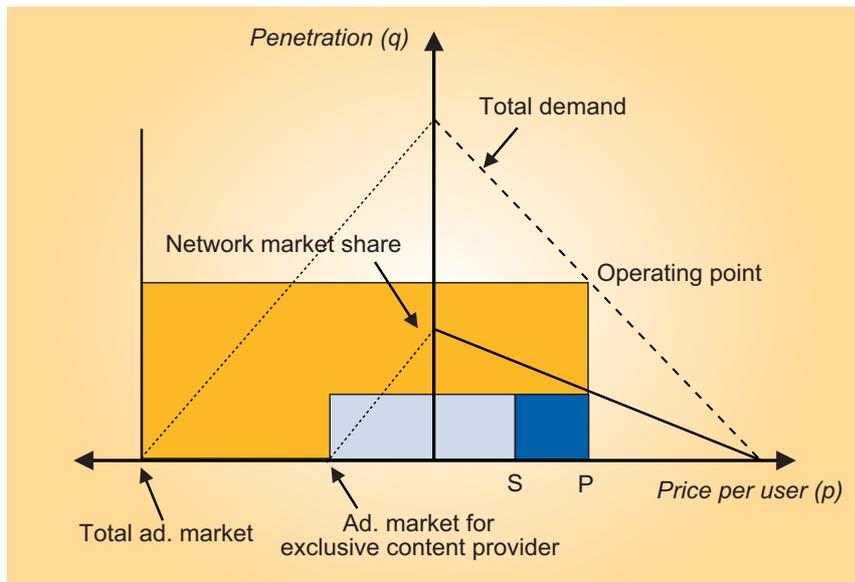


Figure 5
With multiple networks every network must recover its fixed costs, leading to a higher end user price P for all networks (with price competition). Exclusive content deals limit the advertising market, requiring a higher revenue share S from the network and further increasing end user price.

Fig. 6 shows a pure pay-TV context. For now we will only consider ad-free content, paid for entirely by sharing the revenue of the end-user payments. Which of the parties “owns the customer” is to a large degree irrelevant for the analysis, though it can make a difference in the negotiations.

Suppose the network operator receives the user’s payments, and is confronted with content owners demanding their fair share of the revenues. The deals are usually concluded as a per subscriber fee, and when the negotiations have taken place, the network operator is confronted with a total sum (C) to pass on to the content providers. The best operating

point now lies about halfway between C and the end of the demand line. It follows that pay TV can never reach more than 50 per cent penetration: if it does, it is too cheap to be optimal! A simple calculation reveals that the best C for the content providers is halfway along the demand curve, which yields an optimal operating point with a penetration of 25 per cent of the market. The revenue split is then 2/3 for the content providers and 1/3 for the network operator.

This shows an interesting conundrum: the more the content owners demand (on a per user basis) from the network operator, the less they will earn. Increasing the content fee (C) leads to a lower penetration, and consequently less total revenue (the area of the rectangle to the left of C). It would be better for both the network operator and the content providers to base their content fee on the total revenue of the operator, not on the number of subscribers, or make it fixed. This would allow the network operator to lower its price, move the operating point up towards the optimum and increase the revenue. This is however extremely difficult to achieve in practice, because the single content provider who keeps insisting on a per-user fee will win the most. This is a nice instance of the well known “prisoner’s dilemma” from game theory: by each choosing the option they think is best for them, they all end up in the worst possible situation! A way out of this might be total openness of all contracts between the network operator and content providers, so that all parties can verify that the fees are only based on the total revenue (or fixed), not subscriber based. Suppose an agreement can be reached and the operating point rises to halfway along the demand line. Compared with the optimal case above, total revenue would increase by six per cent, and if the content providers would accept the same total sum, the revenue for the network operator doubles.

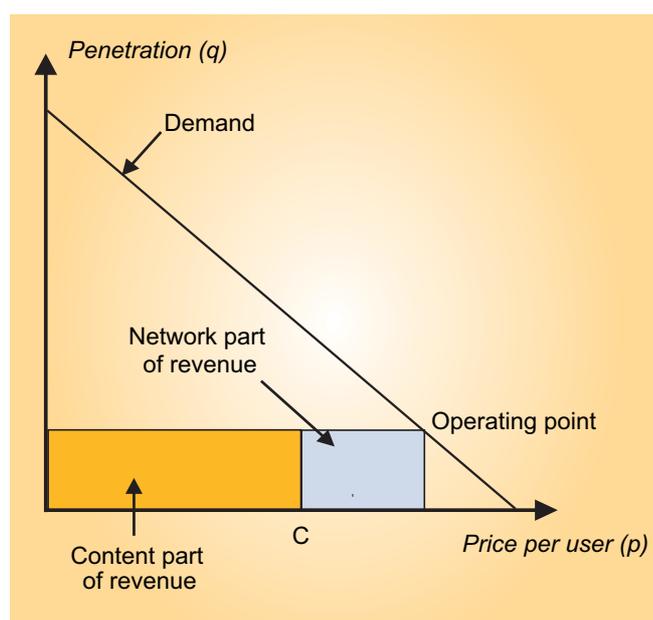


Figure 6
Pay television model

Price differentiation

A well known technique to pull more money out of the market is to address different market segments by selling variants of essentially the same product at different prices, thereby lowering the consumer surplus. In our case this amounts to offering different packages or *bouquets* of channels at different price levels. At first sight this does not seem like price differentiation at all, since we are selling different products. However, operators generally sell a low-cost basic package to all subscribers, and supplement it with optional pay packages. This can then be likened to first-class rail travel, a classical example of price differentiation: all travellers reach the same destination (spend most of their viewing time in the basic package), but some pay more to make the voyage in an exclusive environment (have access to additional bouquets).

When comparing the demand line with only the low-cost utility package (with usually a penetration of over 75 per cent, *Fig. 4*) and the expensive pay TV package (penetration often below 25 per cent, *Fig. 6*), one notices a gap in the middle. Filling this gap requires a package with around 50 per cent penetration. As we have seen in the previous paragraphs, this is sub-optimal both for free-to-air programmes and pay television. However it is quite feasible to assemble such a “low pay” package with channels that are financed jointly by parallel income and by direct payment.

One can imagine an arrangement between content providers and network operators in which the network operator is lured away from its optimum 50 per cent location, by payments from the content providers when penetration exceeds 50 per cent, or, with a similar effect, rebates on the fees paid by the network to content providers. It is easy to see that this revenue sharing is bound by two limits. One limit is where the network revenue is kept constant as the price and penetration move up or down; in the other limit, the content providers keep a constant revenue. In between these limits, negotiations are possible.

The negotiating margin is at its highest at the best “vertically integrated” operating point. With a small parallel revenue of the content providers, the network operator will be dominant, and the penetration will lie close to 50 per cent. Content providers with a high parallel income should demand higher penetration rates from the network operator in return for their payments or rebates. A problem for the content providers here is that the network provider composes the low pay bouquet and will negotiate with every content provider individually. As each of these content providers has only a small parallel income, they will not be able to negotiate a high penetration on their own. It would be best for them to join forces in these talks.

Up to now we have tacitly assumed that the demand line represents the aggregated market of all network operators and content providers. It should be clear that the model also applies to a single content provider. In this case the top of the demand line represents the average market share of this provider. The part to the left of the vertical axis is its parallel income; to the right is the revenue it can obtain through sharing of the network’s revenue. As parallel income per household is taken to be constant and the tail of the demand curve is fixed, the optimum in *Fig. 2* applies to all content providers that tap into the same parallel income, and therefore all have the same optimum position for shared revenue.

Note that in this model, consumer surplus is reduced and turned into additional revenue for the network and content providers. This can help to sustain the operation of multiple networks and relieve the problem we mentioned in that section, without decreasing welfare.

The real demand curve

Up to this point we have assumed a linear demand curve. However it is not difficult to determine the real demand curve. Suppose you ask a representative sample of Germans the following question: “if *Das Erste* were switched off, how much would you be willing to pay in order to get it back?” You would get different answers which, when plotted in a histogram, show the willingness to pay for *Das Erste*. You repeat the question with *ZDF*, *Das Dritte*, *RTL*, *Sat1* etc., all yielding different histograms or probability density distributions. A central theorem in statistics states that when you

bundle these programmes, the resulting distribution of willingness to pay rapidly converges to a normal distribution.

Statistics text books sometimes illustrate this rapid convergence by adding uniformly distributed variables (with a rectangular distribution): adding just three such variables displays the familiar bell-shaped Gaussian curve, while adding six of them produces the normal distribution curve to within a few parts per million! As we are usually dealing with (basic) packages of over six programmes, we

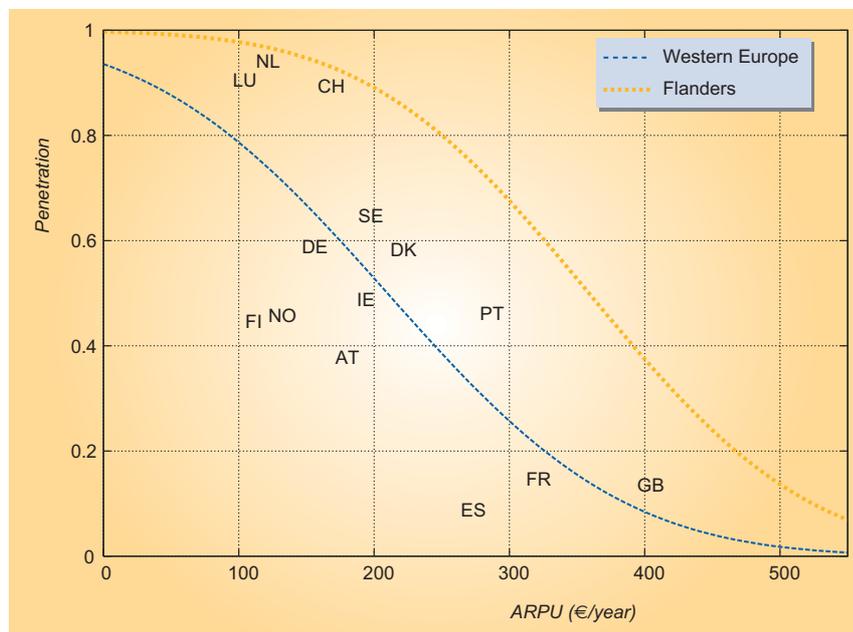


Figure 7
Average demand curve for several European markets and Flanders in 2002, based on figures from Screen Digest, December 2003, p. 364

can safely assume that this is what we were looking for. The curve we need is actually the inverse cumulative distribution, shown in Fig. 7: it shows the number of households willing to take the bundle at a particular price or less. The actual curve shown is the average demand curve for several European countries and the curve for Flanders. Notice how the curves are almost linear in the centre.

As the number of channels increases, so will the average of the sum distribution: the tail of the curve will shift to the right. Beyond a certain number of channels this effect becomes very small, by the so called law of diminishing returns: if you have only one channel to watch, a second one has a high

added value. If you have one hundred to choose from, an additional one will not make much difference.

Generalistic channels address the public at large with a wide range of programmes, each appealing to different segments of the audience. One might therefore assume that the curve even holds for a single generalistic channel, composed of various programmes. At the other extreme, network operators often bundle internet access and telephony with television, the “triple play” offer. According to our theory, this will increase the value, so shift the curve to the right, but conserve its shape.

The normal distribution is completely determined by only two parameters: its *mean* and *variance*, corresponding to the position and the slope of Fig. 7 respectively. This means that we only need two points on the curve to draw it precisely. One of those points is very likely the intersection with the vertical axis (no cost), corresponding to the full market (actually around 95 to 98 per cent of households: most mature markets show a hard core of television *refuseniks*). The second point can usually be determined from the penetration of pay television. As this latter point fluctuates over time, the demand curve is not static but should be reassessed from time to time. Figure 7 was determined from the penetration of basic cable and pay television in 2002. Notice the exceptional right shift of the Flanders curve, typical for markets where the viewers have become accustomed to paying for television, and are mostly unaware of the existence of free to air television!

Changing the linear model above to the real demand curve involves more complicated mathematics, but is not hard for a computer. The results remain largely unchanged, in particular the optima of 50 per cent penetration for a stand-alone network and even the 25 per cent for pay-television. It turns out that this 25 per cent is independent of the market situation (i.e. the mean and variance of the demand curve)!



Kris Van Bruwaene gained an M.Sc. in Electronics Engineering from the University of Ghent, Belgium, in 1975. He obtained a postgraduate diploma in Business Management from KU Leuven, Belgium, in 1998. During ten years in the VRT projects department, he assisted in the introduction of new technologies such as digital video, teletext and electronic subtitling. He then spent ten years in the technical training department, lecturing on digital audio and video signals and devices, digital audio broadcasting (DAB), signal processing, compression and computer processor architectures.

Since 1996, Mr Van Bruwaene has been with the department of Strategy, Technology and Innovation at VRT, with a special interest in the impact of new technology and networks on the future of broadcasting.

Conclusions

We have shown that even a simple model can explain some remarkable properties of the free-to-air and pay television markets. In most cases there are clear optimal operating prices and margins for negotiating the revenue share. The real demand curve can easily be determined. The model reveals that a successful pay-television operator should reach a penetration between 25 and 50 per cent, lowering the price if necessary, and that free to air television is a public good offering the greatest welfare. It also suggests an alternative way of revenue sharing for pay television, based on the total revenue as opposed to the number of subscribers.

Acknowledgement

This work was performed as part of the EU Integrated IST Project *Medianet*.
