A joint Initiative of Ludwig-Maximilians-Universität and Ifo Institute for Economic Research

CESifo

Venice Summer Institute 2001

Workshop on Industrial Organisation

Venice International University, San Servolo
20-21 July 2001

Seeing Double? Duplication, Diversity and the Public Good of Television

Diana Barrowclough

CESifo
Poschingerstr. 5, 81679 Munich, Germany
Phone: +49 (89) 9224-1410 - Fax: +49 (89) 9224-1409
E-mail: office@CESifo.de
Internet: http://www.cesifo.de

DRAFT: NOT FOR QUOTATION OR ATTRIBUTION.

Paper for CESifo Summer Institute 2001, Venice, Italy.

"Seeing double? Duplication, diversity and the public good of television."

Key words: television, advertising, spatial competition, product differentiation.
JEL classification: H,I,L.

Diana Barrowclough
St John's College,
Cambridge,
CB2 1TP.

DVB1000@econ.cam.ac.uk

28 February 2001.

ABSTRACT:

This paper develops an original model of product differentiation, to contribute to the debate about the regulation and finance of public television. It goes beyond the conventional analysis in this topic, by exploring the spill-over effects that a public broadcaster has upon commercial broadcasters. It shows how the existence of a publicly funded, free-to-air, channel (such as the BBC) affects the behaviour of advertiser-financed, free-to-air, channels; including what happens if the public channel becomes less distinctive, and/or introduces advertising.

These are timely issues, given the extent to which public broadcasters are increasingly criticised for seeking popularity, losing distinctiveness, and in many cases, introducing advertising. To illustrate these and other questions of this nature, we develop a general and abstract model, which clarifies the interplay of the issues that are currently causing such tensions in the television industry. (In addition, the model has wider parallels to other sectors where services are offered free at the point of access, but financed by advertising, such as the internet.)

To achieve this, the following pages develop an original model of product differentiation in two dimensions, following the tradition of Hotelling and Cournot competition. The horizontal product attribute is programme quality or type, and the vertical attribute is level of advertising. Broadcasting channels compete for viewers by altering their levels of advertising. The second novelty of this model is its pricing scheme, which captures the unusual nature of television advertising markets. Channels sell quantities of airtime to advertisers, the unit price of which is determined by the number of viewers. Relative demand therefore plays the role of price in a Cournot model, except there can be different prices for different channels.

We use this model to show that there is a trade-off to be made between distinctiveness and advertising. These trade-offs are not always intuitive. Under the assumptions given, we show that if constraining total advertising is the social planner's greatest priority, it is best for the public broadcaster to have zero programme distinctiveness, and zero advertising. This leads to duplication of programme types, but lower advertising. However, if distinctiveness is the top priority, then we must grasp the uncomfortable conclusion that this brings with it higher levels of advertising on the commercial channels. The worst case scenario shown is when the public broadcaster aims to maximise audience numbers and or advertiser revenue: this leads to minimum programme distinctiveness, *and* high levels of advertising.

1. INTRODUCTION[1]:

The most convincing arguments today for publicly financed television are that it provides types of programmes that commercial television would not provide; and that is free of interruptions by advertising. However this traditional "gap filling" rationale has been wearing rather thin lately, as programmes seem to be more or less the same whichever channel one watches; and public channels increasingly either sell advertising directly, or devote large amounts of airtime to self-advertising[2]. Public broadcasters do not seem to be particularly distinctive; and they appear to be competing with, rather than complementing, commercial broadcasters. Perhaps because of this, there is growing public support for abandoning many of the traditions of publicly financed television. Many countries are experimenting with various new forms of finance and delivery of public television, including the establishment of partially competitive quasi-markets, or an increasing use of private sector finance. In the United Kingdom, for example, consumer surveys consistently find about 60% of respondents in favour of financing the BBC through the sale of advertising, instead of the 75-year-old Public Broadcasting Fee[3]. Our question here is, to what extent does this matter?

The importance of "gap filling" arguments notwithstanding, this paper aims to add some fresh insights to the debate about television regulation, by focusing instead on the spill-over effects of public broadcasters. We go beyond the usual analysis, to explicitly include their impact on the behaviour of commercial broadcasters. For example, the existence of a publicly financed broadcaster may affect the amount of advertising transmitted on commercial television. More subtle effects depend on how distinctively different is the public broadcaster's programme profile, or whether it also sells advertising. We show, for example, the danger of introducing even a small amount of advertising onto the public broadcaster, because of the way it encourages both less distinctiveness, and an escalation of advertising in general. These wider effects give an added rationale for the public broadcaster complementing, rather than competing with, commercial broadcasters. Given this context, this paper examines the following questions:

i) How does the existence of a publicly funded, free-to-air, channel (such as the BBC) affect the behaviour of advertiser-financed, free-to-air, channels?
ii) What happens if the public channel becomes less distinctive?
iii) What happens if it introduces advertising?
iv) What is the trade-off between distinctiveness and advertising?

These are timely issues, given the extent to which public broadcasters are increasingly criticised for seeking popularity, losing distinctiveness, and in many cases, introducing advertising. To illustrate these and other questions of this nature, we develop a general and abstract model, which clarifies the interplay of the issues that are currently causing such tensions in the television industry. (In addition, the model has wider parallels to other sectors where services are offered free at the point of access, but financed by advertising, such as the internet.)

---

[1] This paper represents part of my PhD dissertation on broadcasting economics, entitled 'Seeing double? Duplication, diversity and the public good of television.' It has benefited immensely from the encouragement and advice of A.B. Atkinson, and A. Shaked. Any errors remain, however, my own.
[2] Empirical surveys on this include Blumler (1986, 1991), Ishikawa et al (1996), Hillve et al. (1997), and Barrowclough (2000).
[3] 'Attitude survey on behalf of BBC Funding Review', April 1999, MORI, UK.

To answer these questions, the following pages develop an original model of product differentiation in two dimensions, following the tradition of Hotelling and Cournot competition. The horizontal product attribute is programme quality or type, and the vertical attribute is level of advertising. Broadcasting channels compete for viewers by altering their levels of advertising. The second novelty of this model is its pricing scheme, which captures the unusual nature of television advertising markets. Channels sell quantities of airtime to advertisers, the unit price of which is determined by the number of viewers. Relative demand therefore plays the role of price in a Cournot model, except there can be different prices for different channels.

This is an original treatment of the real-world situation with which we are all loosely familiar, in the sense that almost all of us own television sets, with which we (still) mostly watch free-to-air, publicly and/or advertiser financed television. However the issue has not been widely explored in mainstream economics. Certainly, there do not seem to be any other models that look at the problem in this way, with two such attributes.

We use this model to show how the existence of a public broadcaster affects the behaviour of other broadcasters; and then what happens if the public broadcaster starts to become less distinctive, or if it introduces advertising. We show that there is a trade-off to be made between distinctiveness and advertising. These trade-offs are not always intuitive. Under the assumptions given, we show that if constraining total advertising is the social planner's greatest priority, it is best for the public broadcaster to have zero programme distinctiveness, and zero advertising. This leads to duplication of programme types, but lower advertising. However, if distinctiveness is the top priority, then we must grasp the uncomfortable conclusion that this brings with it higher levels of advertising on the commercial channels. The worst case scenario shown is when the public broadcaster aims to maximise audience numbers and or advertiser revenue: this leads to minimum programme distinctiveness, *and* high levels of advertising.

2. THE MODEL.

Assume there are three television channels, delivering three types of programmes, denoted by H, M, and L. (High-brow programmes, Middle-brow and Low-brow). The intuition is that programme type is given, perhaps by a regulator: the High channel is not allowed to provide Low-type programmes, and vice-versa.

Alongside the programmes, each channel can also deliver advertising, denoted by y. The level of advertising can be any amount from 0 (no advertising) to 1 (a lot of advertising): $0 \leq y \leq 1$. Advertising is a source of dis-utility to viewers, so that for any programme type everyone prefers less advertising to more.

2.1 Defining preferences:

Assume that consumers have inherent preferences over each program type. We measure these preferences in terms of the amount of advertising they must absorb on a favoured programme, until the point where they become indifferent between it and a less favoured programme. The intuition is this: if one loves High and hates Low programmes, then utility is only equalised between the two when High comes with a lot of advertising, or Low comes with relatively little or none[4].

[4]The intuition is that of positively-sloped indifference curves. If one plots programme quality or type on the horizontal axis, and quantities of advertising on the vertical axis, quality must rise as advertising rises, in order to maintain constant utility. (Earlier versions of the model treated consumer taste for adverts as a quadratic, with utility rising, peaking and then falling as advertising levels increased, as if consumers understood that advertising

Different consumers have different preferences. Therefore, let the numbers *h,l* describe the preferences of a single consumer, where *h* describes their marginal rate of substitution between High and Medium, and *l* their marginal rate of substitution between Low and Medium. Thus for each level of advertising *y* on Medium, utility is equalised where:

$$U(M, y) = U(L, y + l)$$

$$U(M, y) = U(H, y + h)$$

Where U (X, y) is the utility that the consumer receives from consuming *X*-type programmes, with advertising of the amount *y*. In addition, we assume limits to the preference distribution $-1 \leq l \leq 1; -1 \leq h \leq 1$.

Figure 1 shows this visually. Each vertical line represents the potential range of advertising on programmes Medium and Low, from none at the line's base ($y = 0$), to a lot of advertising, ($y = 1$), at the top. For this person, if Medium is offered with advertising level *y*, utility is equalised between Medium and Low only when Low has advertising levels $y + l$. In this example, the value of *l* is positive, indicating that the person is inherently a Low-lover because they are not indifferent between the two channels until Low is delivered with *l* more advertising than Medium. If both programmes were offered with identical levels of advertising, this consumer would always prefer Low.[5] Figure 2 shows a person whose negative value for *l* indicates that they are by nature a Medium lover: utility between the two channels is equalised only when Low is delivered with much less advertising than the amount of advertising on Medium.

---

was a necessary way of financing the production of television programmes. However little was gained from this added complexity, which was therefore dropped in the following analysis.)

[5] One can imagine an entire map of such indifference curves, stacked in parallel lines above and below the single one shown. When both Medium and Low have the same level of advertising *y*, maximising utility means choosing the *lowest* indifference curve possible, and hence choosing Low. (Note, utility rises as we go down the vertical axis, and not up, as usual, because we always prefer less adverts to more, and $y = 0$ is at the base, while $y = 1$ is at the top.)

Figure 1: Inherent Low-lover's marginal rate of substitution between Medium and Low, where Medium has advertising level *y.*

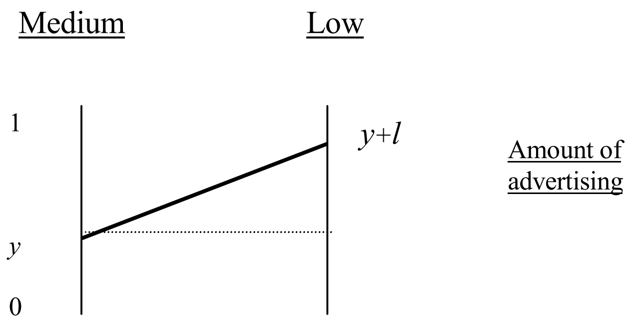

Figure 2: Inherent Medium-lover's marginal rate of substitution between Medium and Low, where Medium has advertising level *y.*

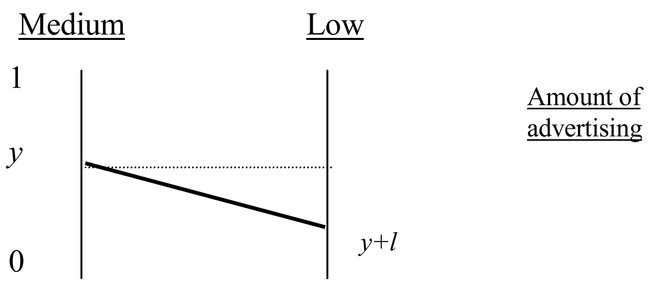

One can imagine conducting this exercise for every person in the population, to graph their values for *l* along the continuum shown in Figure 3 below. Imagine that Medium is provided with advertising levels *y*. Everyone whose value of *l* is positive, given this level of advertising *y*, prefers Low. Those with strongly positive *l* values are located at the far end of the distribution, while those with weakly positive values are located closer to *y*. By comparison, everyone whose *l* is negative, prefers Medium.

Figure 3 : Marginal rates of substitution between Medium and Low, where advertising on M = *y.*

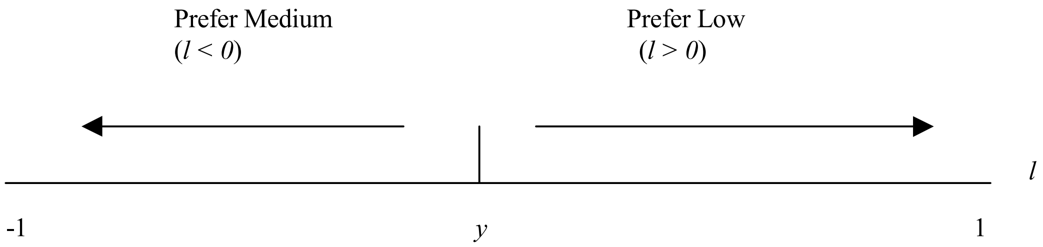

A similar exercise obtains consumers' values for *h* (High compared to Medium, given *y*). These values for *h* and *l* are then distributed in the taste plane. This is drawn below in Figure 4, in the case where no channel actually has any advertising ($y = 0$). We show values for l on the horizontal axis and h on the vertical. (Their limits are not always labelled, to keep the box diagram reasonably simple, but from above, both horizontal and vertical axes are bound between –1 and 1.)

All individuals with taste values $l > 0 > h$ are located in the south-east quarter. They prefer Low to Middle (l>0), and would hence be willing to bear positive amounts of advertisements on Low, when Middle is offered with no advertising, before utility was equalised between the two programmes. Similarly, since 0>h, they prefer Middle to High, meaning they would be willing to bear positive amounts of advertising on Middle, before utility was equalised with High. Any individual located in the north-west quarter has the preference ranking *H*>*M*>*L*, because l<0 and h>0. In the north-east quadrant, both l,h are positive, meaning that both *H* and *L* are preferred to *M*. Splitting that quadrant diagonally distinguishes the ranking *H*>*L*>*M* from *L*>*H*>*M*. A similar diagonal split in the south-west quadrant, where both l,h are negative, separates those people for whom *L* was the worst option from those for whom *H* was worst. Both types preferred *M* best.

Figure 4: Preferences for High, Medium and Low, with zero advertising.

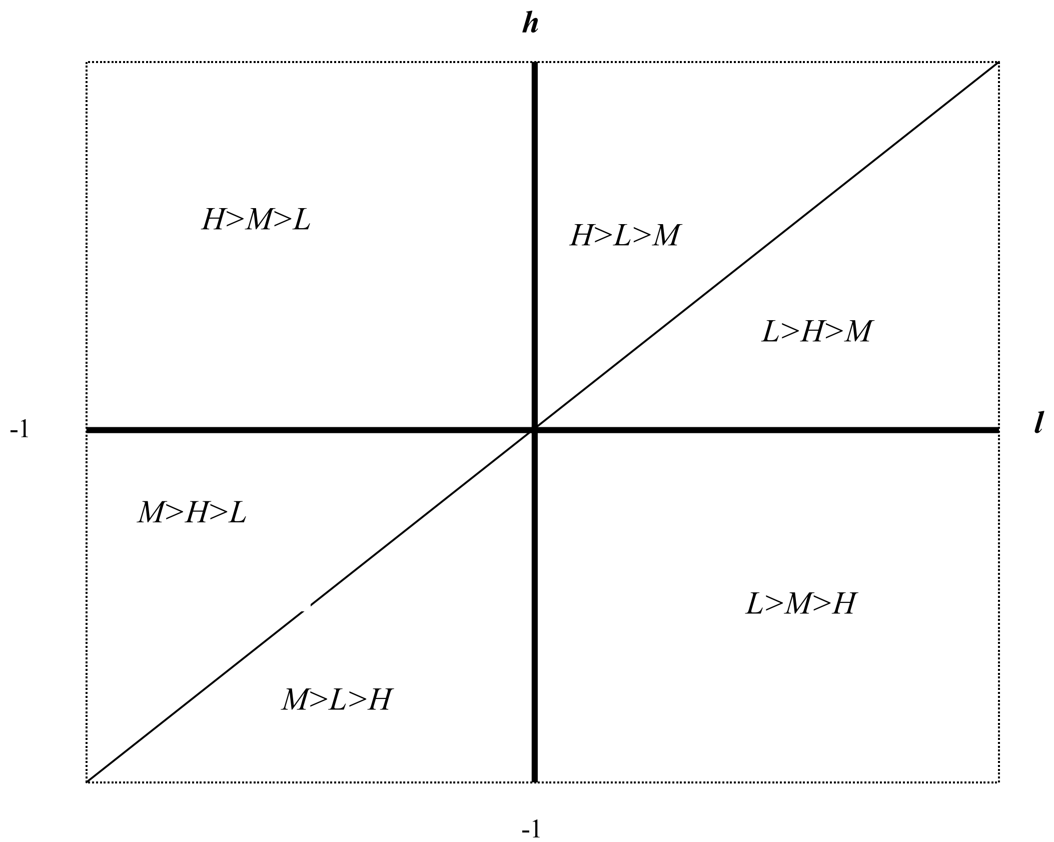

2.2 The effect of advertising.

We can now superimpose onto Figure 4, the effect of differing advertising levels. The three channels and their respective advertising levels are denoted as (H, η), (M, μ), and (L, λ). We also define two variables that depict relative advertising levels: $a = \mu - \eta$, and $b = \lambda - \mu$.

The plane above must now be re-drawn to show individuals' values h,l given the levels of advertising η, μ, and λ, and hence relative advertising values, a and b. Figure 5 shows an example where advertising quantities were greatest on Low, and least on High, so that $\lambda > \mu > \eta$. (As above, both vertical and horizontal axes are bound between –1 and 1, although this is not fully labelled). The higher levels of advertising on Medium compared to Low gives us the new (bold) vertical axis in the south-east quadrant found by $l=b$; the higher levels of advertising on Medium compared to High give us the (bold) horizontal axis in the south-east quadrant formed from the vertical height given by a. This means that the plane of preferences is divided in a new way. Now all individuals with $l>b$ will prefer (L,b) to (M, μ). All individuals with $l > h + \lambda - \mu$ prefer (L,b) to (H, η), and so on.

Comparing Figure 5 with Figure 4 shows that some people who formerly preferred Medium and Low now prefer High; while some viewers who formerly preferred Low now prefer Medium.

Figure 5: Demand for High, Medium and Low, given η, μ and λ.

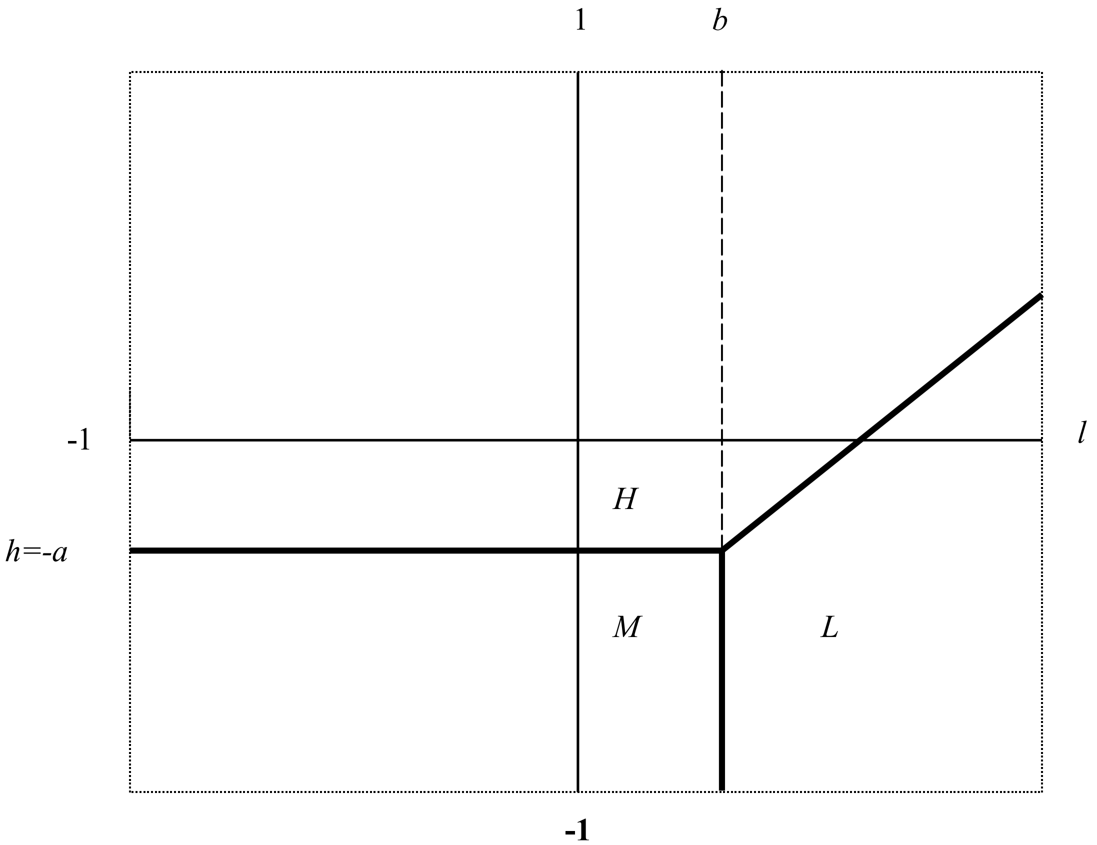

2.3 Assumptions about the distribution of values *h,l.*

Before we can put the model to work, a number of simplifying assumptions are made about the distribution of taste values, and the behaviour of channels. These enable the development of a model that is sufficiently abstract to be able to isolate the television sector's key features in a useful way.

i) We assume that individuals are uniformly distributed within the plane. An alternative assumption could have been that these taste values were bunched or skewed, perhaps reflecting age or income. However it is not appropriate to make such precise assumptions about peoples' tastes, given the abstract nature of this model.[6]

ii) Non-viewing is not an option. This is a hugely simplifying assumption to make, as in reality each individual will have some value of h,l below which they refuse to consume the television offered, but prefer instead to do any other activity (even if only sleeping). Determining the effects of these various "exit" levels would be a fruitful avenue for future research, but the "all view" assumption made here is sufficient for this general and abstract model.[7]

iii) We also restrict the domain of consumer preferences, to the lower rectangle of the plane, given by the dimensions $-1 \leq l \leq 1$, $-1 \leq h \leq 0$, rather than its entire square, given by $-1 \leq l \leq 1$, $-1 \leq h \leq 1$. This is equivalent to saying there are very few people with the preference orderings $H>M>L$; $H>L>M$; or $L>H>M$. Tastes are still distributed uniformly but only within the bottom half of the square shown above. Total density therefore is 2, (ie. 1+1 in each square).

This makes the analysis more tractable, without being too violent to reality. One could argue that it is more plausible to remove the top rather than the bottom half of the plane, as observation suggests that fewer people prefer High-brow to Low-brow than vice-versa.[8] Moreover, of those people who do rank Low as their first preference presumably more prefer the orderly ranking $L>M>H$, (south-east) than the perverse ranking $L>H>M$ (north-east-east).

More importantly though, this simplification reflects the fact most of the "action" in the model below will concern responses made by Medium and Low, with respect to varying strategies chosen by High. Hence we wish to concentrate on the bottom half of the plane. This has the flavour of realism to the extent that the (reasonably unregulated) BBC behaves like High, the (tightly regulated) commercial channel ITV behaves like Medium, and neither offers programmes like Low. (And one need not fear that focusing

[6] This should however be an interesting avenue for future investigation, for example, in the incidence of welfare effects on different social or income groups.

[7] Returning to our imaginary map of indifference curves for each individual (footnote 3), there will exist some pivotal indifference curve beyond which the utility from watching either programme is less than the utility to be gained from all other non-television activities. The location of this pivot point ("watch", "don't watch") is likely to be highly sensitive to age, income, and more generally, the time of day. This would also be a fruitful topic for future research, for those interested in the distributional aspects.

[8] As a digression, one must ask whether High could still mean "Highbrow", if everyone preferred it. For a less emotive titling, Low can be defined as "lowest-common-denominator," rather than lowbrow.

on the southern half of the plane undermines High's strategic arsenal, given the powers allowed to High below.)

iv) The following assumption aims to capture some of the flavour of the current debate in public television, where the BBC is increasingly accused of mimicking the programme range and styles of the ITV. The intuition is as follows. Imagine that Medium and Low are not allowed to change their programme type or intrinsic nature (for example, reflecting the strict regulatory environment in which the advertiser-financed channels operate in the UK). High however has more flexibility, in the sense that its programmes can become more, or less, similar in nature to Medium. (It cannot become like Low.) The degree of High's similarity to Medium will be measured by a "distinctiveness parameter", α, which has the range $0 \leq \alpha \leq 1$. Where $\alpha=0$, High's programme is identical to Medium; where $\alpha=1$, High's programme is distinctively different. Through the use of this distinctiveness parameter, the model is able to illustrate the implications of changing levels of channel distinctiveness and advertising.

High's ability to alter its programme nature is connected to assumption (iii) above as follows. As High becomes more similar in nature to Medium, everyone's values for *h*'s become smaller. (The indifference curves drawn in Figures 1 and 2 become much flatter, because our preferences for programme types are less extreme the more similar they are.) The more High becomes like Medium, the more the distribution of *h*'s becomes concentrated around $h=0$. At the limit when the two channels are identical ($\alpha=0$), all the values for *h* converge to 0. The intuition is clear: why would anyone still have an extreme preference value for *h*, when there is virtually no difference between the two programmes? Who would endure large amounts of advertising on Medium, when High offers exactly the same programme but without the intrusive adverts? To be precise[9], the new distribution is assumed to be homogeneous between $[0\text{-}\alpha]$ where $0 \leq \alpha \leq 1$, and with the density $1/\alpha$.

v) We shall assume that the relative quantities of advertising always preserve the ranking shown in Figure 5 above, where $\lambda > \mu > \eta$. This does not violate what we usually see in television markets; and it is also a useful simplification for the remainder of the analysis below, where we determine channel demands and revenues on the basis of the areas shown geometrically in Figure 5. Any substantial re-structuring of the plane (for example, placing the point of trisection in the lower left hand corner, rather than the lower right hand one) would require re-drawing and re-estimating the demand areas, with little additional insights to be gained.

2.4 From preferences to demand.

Once we know everyone's intrinsic preferences for programme types and their sensitivity to advertisements, we can establish the demands for each channel. This can be calculated

---

[9] One can imagine the base of the southern rectangle squashing upwards towards the horizontal axis. The number of individuals has not changed, and they can still have the same broad (horizontal) range of values for *l* as before. However, rather than being spread within the full depth of the rectangle they are now squashed into a narrower ribbon, that runs along the length of the *l* axis. The depth of this ribbon depends on the value of α, becoming narrower and narrower as α becomes closer to 0. It runs horizontally along *l* but vertically becomes increasingly close towards $h = 0$. When $\alpha = 1$ the "ribbon" expands back to the full depth of the original rectangle. (A rough example would be the way that fluid rises up the sides of a container, when a constant volume of water is transferred from a large container into a smaller one.)

geometrically using the plane of taste values. Consumer demand for each of the three channels, given broadcasters' respective advertising levels ((H,η), (M,μ), (L,λ), and $a=\mu-\eta$, and $b=\lambda-\mu$), is derived from the areas shown in Figure 5 as follows:

$$D_H = \frac{1}{\alpha}\left[(1+b)a + \frac{a^2}{2}\right]$$

$$D_M = (1+b) - \frac{(1+b)a}{\alpha}$$

$$D_L = (1-b) - \frac{a^2}{2\alpha}$$

The intuition for finding these areas is as follows. The depth of the bottom half of the plane is given by α, and the density of the distribution within the plane by $1/\alpha$. Therefore demand for Low, D_L, consists of all those pairs of values (h,l) found in the rectangle with width (on the l axis) of 1-b and depth of 1 (depth α times density $1/\alpha$), minus those located in the little left-hand corner triangle, $a^2/2\alpha$. Similar geometrical methods are used for the other two audiences[10].

2.5 From demand to revenues.

Each broadcaster's total advertising revenue is found by multiplying the amount of advertising sold by its price; as is usual in economics. However a special feature of this model is that the price of each unit of advertising sold is determined by the size of the audience for a particular channel, relative to the total audience. This is a realistic assumption, as television advertisers hope to achieve access to as many consumers as possible and will pay higher prices for access to larger audiences.[11] Each channel's revenue curve is therefore its amount of advertising sold, times its demand.

$$R_H = \frac{\eta}{\alpha}\left[(1+b)a + \frac{a^2}{2}\right]$$

$$R_M = \mu\left[(1+b) - \frac{(1+b)a}{\alpha}\right]$$

---

[10] Here demand for Medium has been found by subtracting High's rectangular demand [(1+b)a. 1/α] from the larger square [1/α.(1+b).α]. One could just as easily have found the area for Medium by multiplying the length (1+b) by height (α -a), over the density 1/α.

[11] This does not occur to the same extent in the print media, where increasing marginal costs mean that advertisers are more selective about their target audiences. In addition, there is an implicit assumption here that advertisers' demand to buy airtime is perfectly elastic, as they will buy all the airtime offered, at the going rate. This is justified by observation within current ranges, and is reinforced by the fact that most countries have found it necessary to impose strict upper limits to levels of television advertising, because of unfulfilled demand at the prevailing prices.

$$R_L = \lambda\left[(1-b) - \frac{a^2}{2\alpha}\right]$$

The first-order conditions for each broadcaster, with respect to their level of advertising, and assuming that each takes the others' behaviour as fixed, are:

$$D_H = \frac{\eta}{\alpha}[1 + a + b]$$

$$D_M = \mu\left[1 + \frac{1}{\alpha} + \frac{b-a}{\alpha}\right]$$

$$D_L = \lambda$$

Given these revenue and first order conditions, we now show the optimal levels of advertising λ, μ, and η, and their implications.

3. PUTTING THE MODEL INTO PRACTICE.

3.1 The general approach:

The model is put into practice below under conditions of partial, and then full, competition.

i) In the partial competition treatment, High sets both its degree of distinctiveness α and its level of advertising η, and Medium and Low must simply respond in the best way they can. They must find their optimal levels of advertising μ and λ given High's first move. Intuitively, this echoes the example where High is a non-competitive public broadcaster, whose behaviour is controlled by a regulator, and where Medium and Low are competitive firms, aiming to maximise revenue within the regulatory environment in which they operate.

   We experiment with different degrees of distinctiveness and advertising on High, to show their effects upon Medium and Low, and their effects on total distinctiveness and advertising. This is therefore a timely illustration of the sorts of issues that are causing increasing tension in broadcasting regulation in the United Kingdom and New Zealand. It shows what happens when High becomes less distinctive, and/or introduces given levels of advertising.

ii) The full competition treatment allows High to choose its optimal level of advertising simultaneously with the other two channels, which also choose their optimal levels. High continues to set its level of α independently (for example, taking it as a given from the regulator), but now it must choose its optimal level of advertising, η, against the best responses of Medium and Low. This situation reflects the example where the public broadcaster can choose its optimal levels of advertisements, within an competitive environment of other commercial broadcasters. It should also offer insights to the UK situation, where the possibility of introducing advertising onto the BBC remains a perennial controversy. The difference between (i) and (ii) therefore is that in the former, advertising levels for High are given, whereas in the latter, High can choose its optimal level.

We do not model the situation where High is allowed to choose its level of programme distinctiveness, α, optimally, for a number of reasons. Firstly, this is not allowed under most country's regulatory regimes, which typically set out numerous parameters within which the public broadcaster must broadly locate itself. High can never choose to be like Low, even if most viewers prefer Low type programmes. Secondly, when High does become less distinctive, it tends to do so incrementally, creeping slighly away from its traditional programme profile, rather than locating in an entirely different section of the taste plane. We show this here by the use of the distinctiveness parameter, α. Finally, of the limited research into television economics that exists, there has been more research already conducted into the issue of how broadcasters choose their programme profile, than in the wider effects of this choice.

3.2 Methodology and the tools of analysis:

Most of the following results are obtained through the Maple mathematical programme, which finds solutions for λ and μ, given different input values of High's distinctiveness and advertising levels, α and η. In the extreme example where High becomes identical to Medium, meaning we have $\alpha = 0$ in the denominator, we find the optimal levels of advertising for Medium and Low by hand, estimating their "best response" curves in the tradition of Cournot competition.

3.3 SUMMARY OF RESULTS:

Key results obtained by the model include the following:

i) The more distinctive is High's programme compared to Medium, the greater the total amounts of advertising on television. The less distinctive is High, the less advertising there is in total. Thus there is the irony that a reduction of one kind of failure increases the other.

ii) There is one important exception to this. In the special case where High's programme is identical to Medium's, we can get a worst-case result where both kinds of failures are maximised simultaneously – we can have minimum distinctiveness, *and* extremely high levels of advertising.

iii) More generally, for any given level of High's distinctiveness, the introduction of advertising on High leads the other channels to further increase their own advertising, causing general escalation of advertising. (The first row of Table 1 below shows this clearly, as High's programme is as distinctively different from Medium as possible, and advertising gradually rises as η rises.)

iv) For any given level of advertising on High, total advertising falls as High's programmes become less distinctive from Medium. This is the general trend of which ii) above is the extreme exception. (This trend is shown most clearly by reading down the first column of Table 1, where advertising on High is zero, and advertising by Medium and Low falls as High's programmes gradually become less and less distinctive.)

v) The Nash equilibrium results reinforce the partial competition results above. Interestingly, they suggest that if High is to be allowed to sell advertising, it may be better to let High determine the level of advertising competitively, while maintaining its level of distinctiveness by regulation. With a medium level of distinctiveness, the Nash level of advertising proved to be rather low, compared to the entire set of possibilities.

4. PARTIAL COMPETITION, MEDIUM AND LOW RESPOND, TO THE CONSTRAINTS SET BY HIGH.

The results summarised above are derived as follows. The first order conditions to the revenue functions described above lead to a cubic equation in *a*, which is solved and plotted as a function of differing values for α and η. From *a* and *b*, we calculate Low and Medium's optimal levels of advertising λ and μ, which equate both Medium and High's demand curves (left-hand-side equation) with their first-order conditions (right-hand side equation). Thus, given α and η, we solve the following equations:

Medium. $$(1+b)\left(1-\frac{a}{\alpha}\right)=(a+\eta)\left(1+\frac{1}{\alpha}+\frac{b-a}{\alpha}\right)$$

Low. $$1-\frac{a^2}{2\alpha}=a+2b+\eta$$

$$a=\mu-\eta$$
$$b=\lambda-\mu$$

Using the values produced we are able to find each channel's demand, revenue, levels of advertising, and price of advertising.

4.1 Results:

In both the tables below, the distinctiveness of High's programmes compared to Medium (α) is shown on the vertical axis and the level of High's advertising (η) is shown on the horizontal. Table 1 shows the total amount of advertising on all three channels ($\eta + \mu + \lambda$), as Medium and Low find their best response to High's (given) choices. Table 2 then shows the breakdown for each particular channel. We see the following:

* Reading the tables horizontally, total advertising increases whenever High introduces and increases advertising, for any given level of distinctiveness. Every increase in adverts on High is matched by further increases on Medium and Low.

* Reading vertically, total advertising falls, as High becomes less distinctive for its given level of advertising.

* The empty cells shown in the bottom right-hand corner of the tables indicate combinations where, given High's distinctiveness and its level of advertising, Medium and/or Low had no possible optimal response (usually they had no audience left).

* The bottom row indicates the special case where High has become identical to Medium. This situation tends to produce unusual results, given the special nature of the strategic responses required of Medium. This case is described at the end of this chapter.

Table 1. Total advertising on all three channels, given High's levels of distinctiveness and advertising (symmetric distribution).

Distinctiveness (α) / Adverts on High (η)	**η = 0**	**η= 0.1**	**η = 0.2**	**η = 0.3**	**η = 0.4**	**η= 0.55**	**η= 1**
α = 1	1.065	1.224	1.378	1.527	1.670	1.876	-
α = 0.7	0.923	1.089	1.251	1.406	1.555	-	-
α = 0.5	0.815	0.987	1.152	1.311	-	-	-
α = 0.3	0.696	0.873	1.041	-	-	-	-
α = 0.1	0.568	-	-	-	-	-	-
α = 0	0.500	0.750	1.000	1.250	1.500	1.750	3.000

Table 2. Levels of advertising on Medium (μ), Low (λ), given High's choices of distinctiveness (α) and advertising (η). (Symmetric distribution)

High's choices of (α) / (η)	**η = 0**	**η= 0.1**	**η = 0.2**	**η = 0.3**	**η = 0.4**	**η= 0.55**	**η= 1.0**
α = 1	0.404, 0.661	0.435, 0.689	0.464, 0.714	0.491, 0.736	0.516, 0.755	0.550, 0.775	-
α = 0.7	0.304, 0.619	0.340, 0.649	0.374, 0.676	0.407, 0.699	0.437, 0.718	-	-
α = 0.5	0.227, 0.588	0.267, 0.620	0.305, 0.647	0.341, 0.670	-	-	-
α = 0.3	0.142, 0.554	0.186, 0.587	0.227, 0.613	-	-	-	-
α = 0.1	0.049, 0.519	-	-	-	-	-	-
α = 0	0.000,	0.100,	0.200,	0.300,	0.400,	0.500	1.000,

	0.500	0.550	0.600	0.650	0.700	0.750	1.000

4.2 Examples – the intuition behind the results.

These general results summarised above are illustrated through three worked examples, to illustrate the intuition behind the method, and the implications of the results. Extreme values are used, to exaggerate the point, in the following scenarios:

(i) Assume High is as distinctively different from Medium as possible, with zero advertising.

(ii) High continues to have no advertising, but becomes more similar to Medium in terms of programme type.

(iii) High introduces advertising, but reverts to its distinctive programme type.

Scenario one:

High is totally different from Medium, so $\alpha = 1$. There is no advertising on High, so that $\eta = 0$. The optimal responses found by Maple (Table 2) are for Medium to have advertising levels of $\mu = 0.404$ and Low to have $\lambda = 0.661$. Medium has the biggest audience (0.749), and earns total revenue of 0.303; Low's audience is smaller (0.661), but because it has a lot of advertising it earns more revenue (0.437).

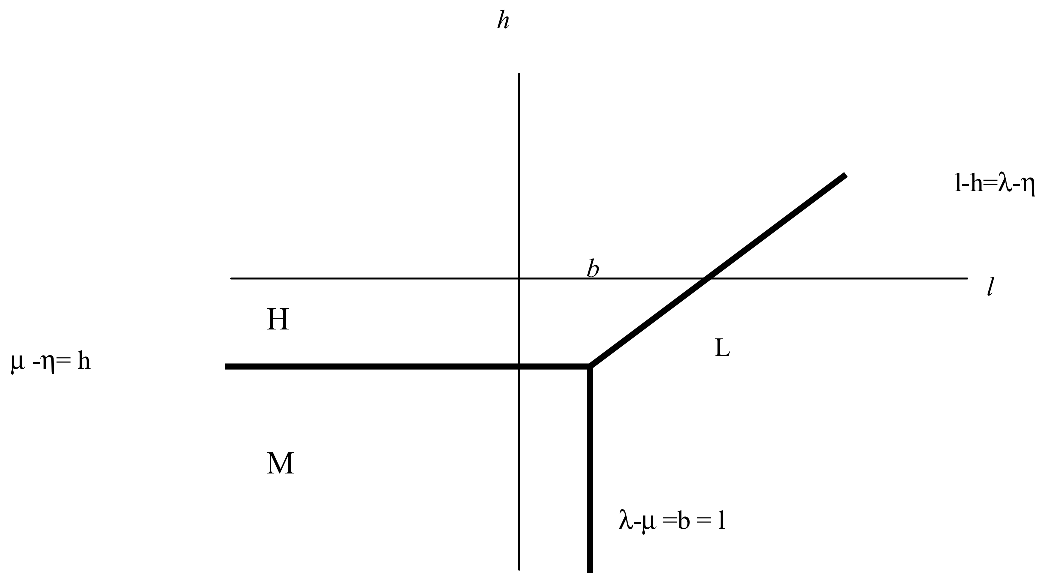

Where: $\alpha = 1$, $\eta=0$, $\mu=0.4$, $\lambda=0.66$, a = 0.4, b = 0.26, H = 0.589, M = 0.749, L = 0.661.

Scenario two:

Now let High become less distinctive, to the extent that it is almost the same as Medium, $\alpha = 0.1$. (The case where High is absolutely identical to Medium is treated separately below). The distribution has therefore become sharply squeezed into the smaller space given by the southern quadrant's shortened depth: all the *h* values are closely located to *h* = 0, although they continue to stretch along the continuum of *l* values. (We show this by shortening the vertical axis in the diagram). Continue to allow no advertising on High, so that $\eta = 0$. Now the optimal responses for Medium and Low are to reduce their amount of advertising: dramatically so in the case of Medium (note the very small value for a), and slightly so in the case of Low.

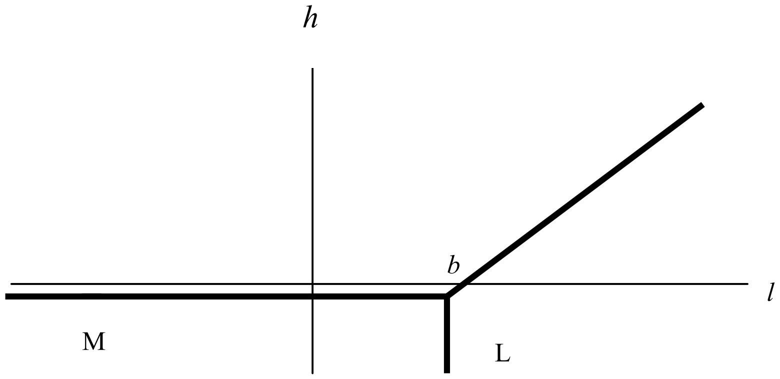

Where: $\alpha = 0.1, \eta = 0$, $\mu = 0.05$, $\lambda = 0.52$, a = 0.05, b = 0.47, H = 0.734, M = 0.747, L = 0.519.

This is intuitively obvious, but it is nonetheless gratifying to find it exhibited so clearly in the model. Medium's former viewers are no longer prepared to endure the same higher levels of advertising as before, given that High's programme is now so like Medium, and moreover has the advantage of coming without advertisements. And as Medium drops its advertising, in order to stop its viewers fleeing to High, so too must Low, in order to stop swing viewers fleeing from Low into Medium. Therefore Medium and Low reduce advertising to $\mu = 0.05$ and $\lambda = 0.52$, preserving audiences of 0.747 and 0.519 respectively. High still manages to increase its share of viewers, at the expense of the other two. Low is not affected to the same extent as Medium, who has High breathing down its neck, but it still experiences a decline in viewer numbers and almost a halving of revenue (from 0.437 to 0.268.) Medium's audience has not changed so much, given that its very small amount of advertising has stemmed the flow of viewers into High. However this small quantity of advertising, despite its high price, brings a much smaller revenue (0.037). (One must wonder whether, in reality, Medium could afford to stay in the market.)

Interpretation: Distinctiveness has fallen, as there are now really only two types of programmes from which to choose: Low, and then either Medium or High, which are more or less the same except that one has adverts and the other does not. Viewers have less choices than before, and both Low and Medium have lower revenues than before. If however the regulator was more concerned about excessive advertising, than excessive "sameness", it will be encouraged to see total advertising substantially reduced, from 1.065 to 0.568.

Scenario three:

Assume High is once again as different from Medium as possible, with $\alpha = 1$. Assume it is allowed advertising, to the level $\eta = 0.55$. Under these parameters, High's share of viewers falls sharply, reflecting the fact that there are very few people who find that they gain greater utility from watching High with this much advertising, rather than watching the very different programmes on Medium, with their advertisements. (This was in fact the maximum level of advertising possible on High, as anything above this caused High to lose its viewers totally, as evidenced by the empty cells above $\eta = 0.55$.)

Therefore it is no surprise to find that Medium can increase advertising to match High, $\mu = 0.550$; and still attract more viewers (1.2243). Hence its advertising revenue increases to 0.674.

In addition, Low can increase advertising to 0.775, attract more viewers, and increase total revenue to 0.6009. By comparison, High now has an extremely small audience of only 0.00049, comprising those individuals located along the horizontal axis, where $0 < h < b$, and making and revenue of 0.00027. As above, in the real world this would seem highly unsustainable.

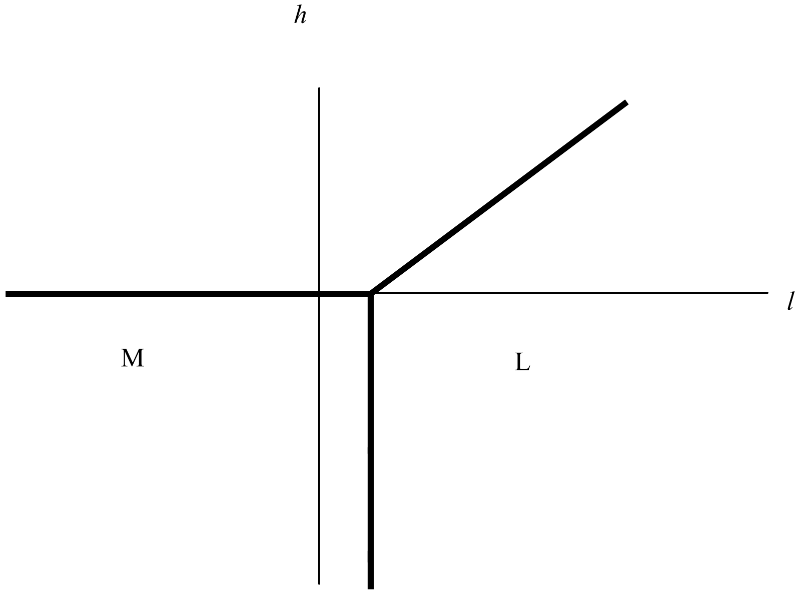

Where: $\alpha = 1$, $\eta = 0.55$, $\mu = 0.55$, $\lambda = 0.775$, a = 0, b = 0.225, H = 0.00049, M = 1.2243, L = 0.775.

Interpretation: This result shows that High suffers dramatically if it tries to remain distinctive in its programme type, and also show high levels of advertising. This does of course stem from our initial assumption that there are relatively few people who have High as their first preference, but the result should nonetheless offer interesting insights to those critics who call for the introduction of advertising on the BBC. The model shows firstly that if selling advertising is an aim, then High needs to reduce its distinctiveness (advertiser revenue in this model was minimal as the high level of advertising was spread over such a small audience). This may not be the result that the regulator wants. Secondly, for any given level of High's distinctiveness, any increase in advertising by High is matched by increases by Medium and Low. Again, this general escalation of advertising is probably not what the regulator had in mind.

5. FULL COMPETITION BETWEEN ALL THREE FIRMS – THE NASH EQUILIBRIUM.

Now we assume that High is able to set its levels of advertising competitively, finding its "best response" alongside the competitive decisions made by Medium and Low. We continue to assume however that High's degree of distinctiveness is determined by the regulator, and hence given. We therefore add a third equation to the Maple programme, to find the value of η that equates High's demand curve with its first-order-condition:

High. $$(1+b)a+\frac{a^2}{2}=\eta(1+a+b)$$

Medium. $$(1+b)\left(1-\frac{a}{\alpha}\right)=(a+\eta)\left(1+\frac{1}{\alpha}+\frac{b-a}{\alpha}\right)$$

Low. $$1-\frac{a^2}{2\alpha}=a+2b+\eta$$

$$a=\mu-\eta$$
$$b=\lambda-\mu$$

As supplied by Maple, the equations above again lead to a cubic equation in *a*, from which we can find the optimal values of η,μ, and λ, for any given α. (The solution to this is shown in the Annex).

The results reinforce those gathered in the two-firm competition case above.

- As High's distinctiveness diminishes, there will be less advertising in total. High gets larger audiences, at the expense of Low's audience. (Medium retains its audience by sharply dropping its advertising). As distinctiveness increases, there is more advertising in total, and High gets smaller audiences.

- Generally, revenue for all channels, including High, is maximised when High is very distinctive. High is therefore generally better off with a high degree of distinctiveness and (compared to Medium and Low) relatively small amount of advertising. This also offers consumers the greatest degree of choice of programme types, and the least advertising.

- The exception to this is when High is allowed to become completely identical to Medium. In this case, everyone's revenues are maximised at very high levels of advertising. This result means that consumers have the least degree of choice of programme types, and the most advertising.

- These results imply that if High is "allowed" by the regulator to duplicate Medium, it has every incentive to do so. If on the other hand it is required by the regulator to be distinctive, it must also be "allowed" to survive on very small audiences, and relatively low advertiser revenue.

Table 3: Full competition - Nash equilibrium levels of advertising, given High's level of distinctiveness (symmetrical distribution).

Distinctiveness	Eta	Mu	Lamda	Total adverts	Audience High	Audience Medium	Audience Low
$\alpha = 1$							
$\alpha = 0.8$	0.1943	0.4034	0.6744	1.2721	0.3595	0.9388	0.7017
$\alpha = 0.7$	0.1770	0.3661	0.6575	1.2006	0.3744	0.9426	0.6830
$\alpha = 0.5$	0.1371	0.2812	0.6198	1.0381	0.4065	0.9529	0.6406
$\alpha = 0.3$	0.0891	0.1811	0.5764	0.8466	0.4418	0.9677	0.5905
$\alpha = 0.1$	0.0321	0.0646	0.5270	0.6236	0.4799	0.9878	0.5323
$\alpha = 0.0$	1.0000	1.0000	1.0000	3.0000	0.5000	0.5000	1.0000

Table 4: Full-competition – Nash equilibrium revenue, given High's distinctiveness, at optimal advertising levels. (symmetrical distribution).

Distinctiveness	Rev H	Rev M	Rev L	Total revenue
$\alpha = 1$				
$\alpha = 0.8$	0.0699	0.3787	0.4732	0.9218
$\alpha = 0.7$	0.0663	0.3451	0.4491	0.8604
$\alpha = 0.5$	0.0557	0.2679	0.3971	0.7208
$\alpha = 0.3$	0.0394	0.1752	0.3404	0.5550
$\alpha = 0.1$	0.0154	0.0638	0.2805	0.3597
$\alpha = 0.0$	0.5000	0.5000	1.0000	2.0000

6. OBTAINING RESULTS WHEN HIGH IS IDENTICAL TO MEDIUM ($\alpha = 0$). THE CASE OF "ZERO DISTINCTIVENESS".

In the special case where $\alpha = 0$, High and Medium must share exactly the same, tightly-packed, distribution of preference values. We begin the analysis by showing what the channels Medium and Low would do in the two-firm case, in the absence of High, before introducing High back into the picture, to show the effect of its presence as a duplicator of Medium-type programmes.

As $\alpha = 0$, the distribution is now no longer within the plane, as in earlier figures, but now lies concentrated along the horizontal line $-1<l<1$ (ie. for every pair of values h,l, the value for h is 0.) The distribution therefore looks like Figure 6 below, with density $1/\alpha$, and a population of 1 in each half of the horizontal axis, and 2 in total:

Figure 7: Distribution of h,l of density $1/\alpha$, where $\alpha = 0$.

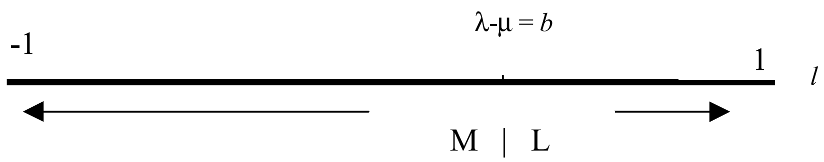

The revenue function for Low continues to be given by D_L times its quantity of advertising: $R_L = (1 - b)\lambda$, or $(1-\lambda + \mu)\lambda$. Deriving this with respect to λ we obtain:

$\delta R_L/\delta\lambda = 1 + \mu - 2\lambda = 0$

So that Low's optimal level of λ, for any μ chosen by Medium is:

$\lambda^* = (1 + \mu)/2$.

Medium's best responses are obtained in the same way. From Figure 7, Medium's revenue is given by: $R_M = (1 + \lambda - \mu)\mu$. Deriving this with respect to μ, we obtain:

$\delta R_M/\delta\mu = 1 + \lambda - 2\mu = 0$

So that Medium's optimal level of advertising for any level chosen by Low, is where:

$\mu^* = (1 + \lambda)/2$

These two "best responses" are plotted on the graph below, as in the familiar presentation of Cournot-type models, with which this shares some similarities. The horizontal axis measures Low's levels of advertising λ, ranging from 0 to 1, and the vertical axis shows values for Medium's μ, also from 0 to 1. The upper curve depicts Medium's best responses to Low's choices of any λ, while the lower curve shows Low's best responses to Medium's choices of any μ. (when Medium chooses $\mu = 0$, Low's best response is to set $\lambda = ½$, etc.)

Figure 8: Best responses for Medium and Low.

Medium's best μ given Low's choice of λ

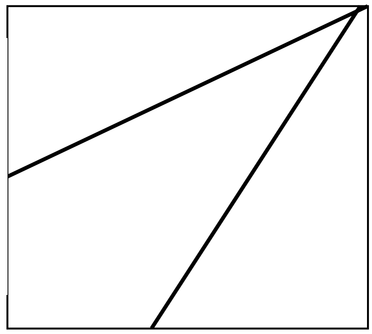

0 ½ 1

Low's best λ given Medium's choice of μ

In this example, the equilibrium response for both broadcasters is shown where $\mu^*=\lambda^*=1$. Revenues are maximised for both when advertising is equal, at its maximum, and the market evenly shared. This is not entirely surprising, given that there are only two broadcasters, and we did not allow non-viewing. The interesting part is what happens next, when we incorporate High into the scenario.

If High enters the market with a programme identical to Medium's, Medium has an extremely limited range of responses. Clearly Medium can never have any more adverts than does High, because if it did it would lose all its viewers. Why would anyone watch a programme with disutility-causing advertisements on it, if an identical programme is offered elsewhere without them? If High has zero advertising, so too must Medium. Or if High has some positive level of advertising, Medium must match it. Moreover, should Medium try to undercut its level of adverts to somewhere just below High's, then High may respond in turn cutting its level again. There are further ripple effects, because as Medium reduces its levels of advertising, then so too must Low. This tit-for-tat behaviour simply results in lower revenues for everyone, Low included, because every time Medium reduces its advertising, Low faces the risk of losing viewers. Medium is best off by responding with exactly the same amount of advertising that High chooses. As shown on the table below, we see that if, for example, High chooses $\eta = 0.5$, μ will be 0.5, and Low reduces its advertising to set $\lambda^* = (1+\mu)/2 = 0.75$.

Table 5: Advertising, demand and revenues when High is identical to Medium – Symmetrical demand.

Eta	Mu*	Lamda*	Audience High	Audience Medium	Audience Low	Rev. High	Rev. Med.	Rev. Low	Total adverts
0.00	0.00	0.50	0.75	0.75	0.50	0.00	0.00	0.25	0.50
0.10	0.10	0.55	0.73	0.73	0.55	0.07	0.07	0.30	0.75
0.20	0.20	0.60	0.70	0.70	0.60	0.14	0.14	0.36	1.00
0.30	0.30	0.65	0.68	0.68	0.65	0.20	0.20	0.42	1.25
0.40	0.40	0.70	0.65	0.65	0.70	0.26	0.26	0.49	1.50
0.50	0.50	0.75	0.63	0.63	0.75	0.31	0.31	0.56	1.75
0.55	0.55	0.78	0.61	0.61	0.78	0.34	0.34	0.60	1.88
0.70	0.70	0.85	0.58	0.58	0.85	0.40	0.40	0.72	2.25
0.80	0.80	0.90	0.55	0.55	0.90	0.44	0.44	0.81	2.50
0.90	0.90	0.95	0.53	0.53	0.95	0.47	0.47	0.90	2.75
1.00	1.00	1.00	0.50	0.50	1.00	0.50	0.50	1.00	3.00

7. SUMMARY AND CONCLUSION.

We have developed above a model of product differentiation in two dimensions, to show the effects of broadcaster decisions with respect to programme distinctiveness, and the level of advertising. This is a spatial competition model that follows the tradition of Hotelling and Cournot: television chanels differ in terms of their programme type and they compete for viewers by varying their levels of advertising. Relative demand (the relative size of each channel's audience) plays the role of price in a Cournot type model, except that there can be different prices for each of the three channels. Hence we have combined elements relating to viewer preferences, their sensitivity to advertising, the broadcasters' chosen level of advertising, and the distinctiveness of the public broadcaster's programme type.

Using this model we are able to show that the existence of a public broadcaster not only can have direct effects, in terms of its programme choices; but that it has spillover effects onto the commercial broadcasters that share its environment. This illustrated some of the trade-offs that have to be made between distinctiveness and advertising. For example, given the assumptions made about the distribution of tastes, we found that when the public broadcaster had a very distinctive programme profile, and zero advertising, consumers had the benefit of being offered three different types of programmes, but with a high amount of advertising on two of them. This occurred because Medium and Low could "get away with" high amounts of advertising, when very few people preferred High type programmes.

Secondly, if High's programme profile changed to become more like Medium's, but without advertising, viewers suffered in the sense that they lost diversity of programme choices, but they also benefited as total advertising levels fell. The final effect depends upon which failure we think is more important: programme choices, or advertising. Thirdly, if the public broadcaster has a very distinctive programme type, but also sells advertisements, there is an escalation of advertising in total. Consumers do get the benefit of three distinctively different programme types, but there is more advertising than occurred in the first instance. And finally, if the public broadcaster becomes identical to another broadcaster as well as having advertisements, consumers have only two different programme types from which to choose, and high levels of advertising on all three channels. This appears to be a worst case scenario.

This paper therefore offers a new rationale for ensuring that public broadcasters complement, rather than compete with, their commercial counterparts. It also shows the danger of introducing even a little advertising onto the public channel. The biggest problem however is that, by definition, a highly distinctive programme profile is also often one that is unpopular, in the sense that it never captures the large audiences that would be attracted to other, less distinctive programme types. The small audience may value it very highly, but it is viewer numbers, and not the intensity of their preferences, that is measured. Most public broadcasters have been unwilling to deal with the tensions that this creates: how does one reconcile the use of public finances to provide a service that very few people consume? Most public broadcasters see their mission as a tight-rope act, of providing programmes to as many tax-payers as possible, whilst at the same time being distinctive and high quality. For example, we are told that public broadcasting in the UK should "provide something for everybody, making the good popular, and the popular good[12]". This inability to resolve the conundrum between distinctiveness and unpopularity is at the heart of many broadcasters' incremental slide towards programme duplication. It is further exacerbated if public broadcasters are also required to supplement their budgets through the sale of advertising.

Ends.

[12] Secretary of State for Culture, Media and Sport, speech to the Royal Television Society, 14.10.1998.

Selected Bibliography.

Arthur Anderson (1998), `The impact of digital television on the supply of programmes', A report for the European Broadcasting Union.

D. Barrowclough (1998), `Public television quasi-markets in New Zealand and the United Kingdom,' in W. Bartlett et al, `A Revolution in Social Policy: Quasi-Market Reforms in the 1990s', The Policy Press, Bristol, pp. 257 – 275.

T.C. Bergstrom and R.P. Goodman (1973), `Private demands for public goods', American Economic Review, June, 63, pp. 280-296.

Bewley (1981), `A Critique of Tiebout's Theory of Local Public Expenditure', Econometrica,49 (3), p.p 713-740.

J.G. Blumler, (1991), `In pursuit of programme range and quality', Studies in Broadcasting, 27, Japan Broadcasting Corporation, pp 191-206.

J.G. Blumler, M. Brynin, T.J. Nossiter (1986), `Broadcasting Finance and Programme Quality: an international review', European Journal of Communication, vol 1, pp. 343-64; and in West Yorkshire Media in Politics Group, 'Research on the range and quality of broadcasting services'; a report for the Committee on Financing the BBC, London:HMSO.

A. Briggs (1979) `The History of Broadcasting in the United Kingdom,' Oxford University Press.

E.H. Chamberlain, (1950), `Product heterogeneity and public policy', American Economic Review, pp. 85-92.

R.H. Coase (1947), `The origin of the monopoly in broadcasting in Great Britain', Economica.
….(1950), `British Broadcasting: A study in monopoly', Longmans, UK;
….(1961), `Why not use the price system in the broadcasting industry', The Freeman, July, pp. 52-57
....(1966) `The Economics of Broadcasting and Government Policy', American Economic Review, 56 (2), May pp. 440-447.
....(1974), `The Lighthouse in Economics', Journal of Law and Economics, 17, pp 357-76.

W. S. Comanor and T. A. Wilson (1979), `The Effect of Advertising on Competition: A Survey', Journal of Economic Literature, 17 (2), June, 453-476.

T. Congdon et al (1992) `Paying for Television', Routledge, UK.

J. G. Cullis and P. R Jones (1998), `Public Finance and Public Choice,' Second Edition, Oxford University Press, Oxford, United Kingdom.

G. Davies (1999), 'The future funding of the BBC', Report of the Independent Review Panel, Department for Culture Media and Sport, London.

P. Hillve, P. Majanen and K. E. Rosengren (1997), `Aspects of quality in TV programming. Structural diversity compared over time and space', European Journal of Communication, Vol. 12 (3), pp. 291-318

H. Hotelling (1929), `Stability in competition', Economic Journal, 39, pp. 41-57.

S. Ishikawa (1996), 'Quality assessment of television', University of Luton Press, UK.

T.Nilssen (1997), `Sequential location when transportation costs are asymmetric', Economics Letters, 54, pp 191-201.

T. Nilssen and L. Sorgard (1998), `Time schedule and programme profile; TV news in Norway and Denmark', Journal of Economics and Management Strategy; 7(2) Summer 1998, pp 209-235.

R. Noll, M. Peck, J. McGowan (1973), `Economic Aspects of Television Regulation', The Brookings Institution, USA.

Veljanovski (1992), `Regulators and the Market, An Assessment of the Growth of Regulation in the UK', IEA, UK
....(1989) `Freedom in Broadcasting', IEA, UK.

Vogel, H.L. (1990) 'Entertainment industry economics: a guide for financial analysis' 2nd ed. Cambridge: Cambridge University Press