Learning from ads: an economic approach to TV content curation

Joaquín M López Muñoz, Valia Merino Vallina, Iñaki Martínez-Sarriegui

Optiva Media, Spain
{joaquin.lopez, valia.merino, inaki.martinez}@optivamedia.com

Abstract

TV content curation has become crucial as Pay TV aggregated offers grow in terms of size and diversity of sources and modes of watching (linear, catch-up, SVoD). Somewhat surprisingly, though, curation is still regarded as mostly an art, dependent on the expertise and intuition of editorial teams. What little formal support there is focuses on the impact of curation (A/B testing, performance KPIs) rather than the curation activity itself (identification of targets and editorial venues, content selection, commercial planning).

On the other hand, TV advertising is a notoriously well understood market with its own terminology, analysis tools and normalized procedures. The backbone of TV advertising is a multi-agent value chain revolving around key economic elements (inventory, rating points, CPMs) with utility values that can be measured and traded on.

We present some ideas on how the key economic concepts of TV advertising can be translated and adapted to the value chain of TV curation so as to provide a formal framework that helps editorial teams guide and assess their activity. These results are intended to spur a conversation within the Pay TV industry on formalizing and optimizing the currently onerous and unpredictable activity of content curation.

Introduction

The uncertain status of TV content curation

Pay TV services have evolved from basic aggregation of linear channels to complex end-user offers combining multiple linear and non-linear sources (channels, TVoD, external and internal SVoD) with different live and on-demand watching modes. The **active window** (content entries available to the end-user at any given time) of medium-sized TV operators can now range in 100,000-150,000 references. In order to help users to navigate these massive catalogues, Pay TV operators rely on several **discovery mechanisms** such as content categorization, automated personal recommendations and **content curation**.

Of all TV discovery mechanisms curation is the least automated and most onerous, and basically relies on the work of human editorial teams in charge of identifying and grouping potentially relevant content and making it available at selected locations within the TV service UX. Curation effectiveness can be measured via automated tools such as A/B testing or user activity/churn KPIs, but these are one-way assessment techniques that do not help editorial teams on the very process of curation nor provide any clue on effectiveness improvement: curation is then regarded as a kind of art, rather than a well-established, reasonably predictable activity.

The economic underpinnings of TV advertising

By contrast, TV advertising is a mature business based on a solid and well understood **value chain** that has been in operation for decades. This value chain can be modelled as a supply-demand market where the supply side (TV channels) trade their **inventory** (slots for ad placement) against the demand side (advertisers) based on an economic valuation of the inventory (classically, **rating points** or **RPs**) reflecting the inventory effectiveness in terms of reach and advertising impact. The model can be further complicated with the introduction of concepts coming from the online world such as alternative valuation metrics (like **cost per mile** or **CPM**), audience segmentation (**targeted advertising**) and automated **inventory brokering**.

The fact that TV advertising runs on economic, measurable principles makes it amenable to formal analysis and automated assistance for decision taking and effectiveness assessment (**ROI**). We would like to have the same level of support in the tricky business of TV curation.

In what follows, we distil the essential economic principles of TV advertising and map them to construct a formal model of TV curation similarly guided by concepts of supply and demand, economic valuation and return maximization.

Basic economic concepts of TV advertising

At its very simplest, TV advertising reflects the economic relationships between two types of agents:

- **Advertisers** who want to promote their products through ads.
- **Publishers** (TV channels, Pay TV operators, OTT players) with reserved ad space or **inventory** available for ad insertion at a price.

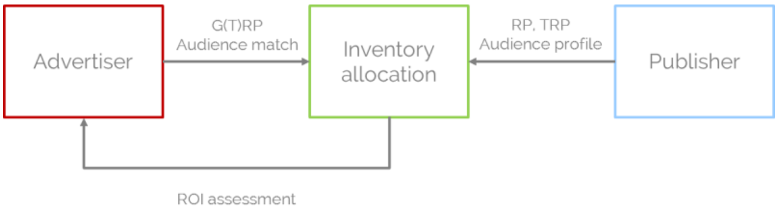

Fig 1. The basic economic model of TV advertising.

The process of agreeing on the placement of a particular ad into an inventory slot is called **inventory allocation**. The publisher qualifies its audience (or has it qualified by an external company such as Nielsen) by numerical means such as **rating points** (**RPs**, percentage of the total audience reached) and **targeted rating points** (**TRPs**, relative to a given audience segment), as well as by non-numerical **audience profile** descriptions. The advertiser plans their ad campaign based on impact estimations (typically measured in **gross rating points** or **GPRs**) and an assessment of the level of correlation or **match** between the audience of the publisher and their customer target population. All these factors play a role in fixing the price for ad insertion. The **return of investment** (**ROI**) for the ad campaign is subsequently assessed by the advertiser as a function of how product sales, brand awareness, etc., have increased.

The negotiation process between advertiser and publisher can be done on a per case basis or, as is the case for online advertising, be automated in a value chain with many agents and intermediaries on either side.

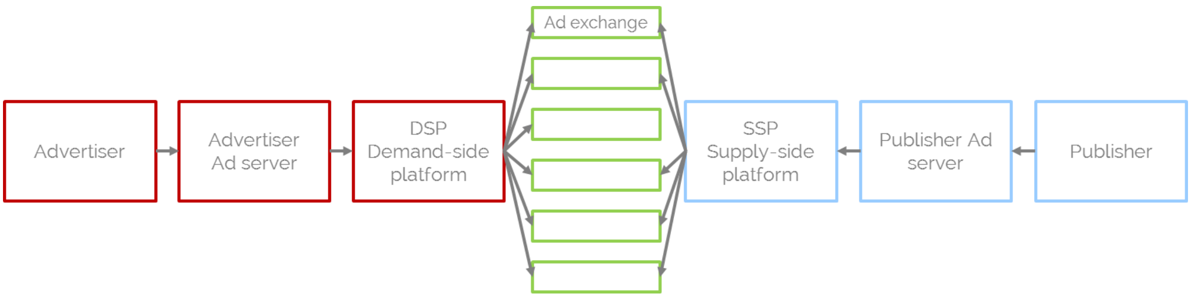

Fig 2. Addressable/online scaleup/automation of the advertising value chain.

Be it automated or conducted in a more traditional fashion, TV advertising is ultimately a very typical supply-demand market.

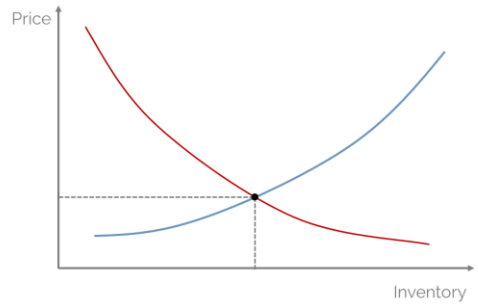

Fig 3. TV advertising as a supply-demand market.

Like other markets of these kind, we observe saturation effects both on the side of the demand (ROIs stagnate from a certain level of ad exposure) and the supply (publishers tend to have lower rating points as their TV services are overcrowded with ads).

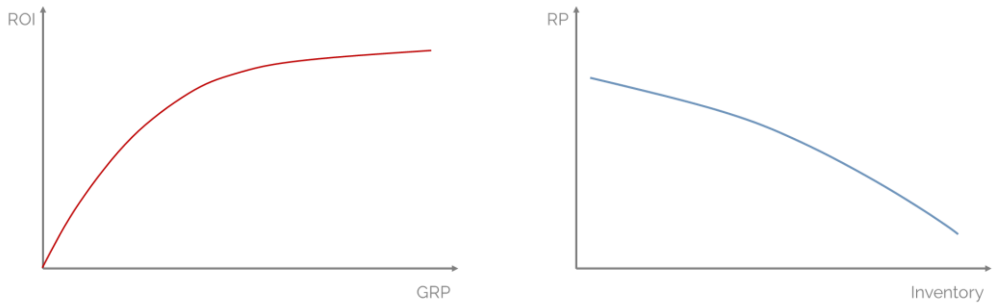

Fig 4. Demand-side and supply-side saturation in TV advertising.

We have mentioned RP and TRP as customary indicators for ad impact estimation and inventory valuation. Table 1 summarizes other typical indicators for specialized scenarios with addressable TV advertising capabilities and/or on-demand content playback. As the market evolves to full-digital, valuation mechanisms will tend towards pricing based on exact impression/impact count (**cost per mile** or **CPM**, **cost per click** or **CPC**).

	Non-addressable	Addressable
Live TV	RP	TRP
On-demand TV	C3/C7	CPM/CPC

Table 1. Inventory valuation mechanisms for different TV advertising scenarios.

An economic model for TV curation

We study the prototypical case of a TV service with a dedicated editorial team in charge of selecting which content entries from the current catalogue ought to be promoted in dedicated **curation slots** within the service UX (for instance, a "recommended" section on the homepage, curated thematic playlists, etc.)

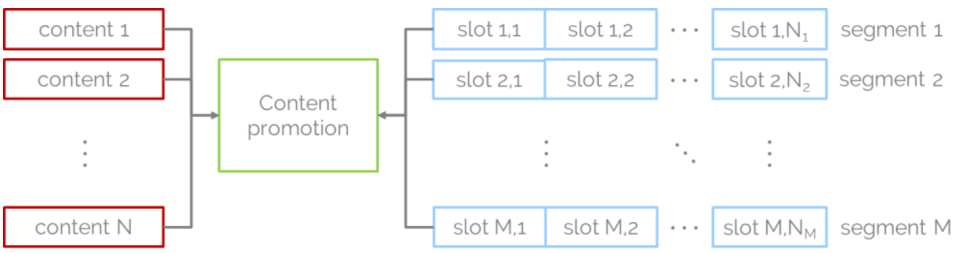

Fig 5. The basic model of TV curation.

We also include a notion of audience **segments**, typically modelled after the information provided by the service BI, which make it possible to do different kinds of curation to different targets (based on age, socioeconomic status, etc.)

To keep things simple, we consider both the segments and their associated curation slots to be fixed —that is, the editorial team works with a predefined curation space that doesn't change frequently.

The key insight to our approach is that this model is, to a large extent, formally equivalent to that of TV advertising if only we do a translation of vocabulary between both domains.

TV advertising	TV curation
Ad	Curated content
Inventory	Curation slots
Inventory allocation	Content promotion
TRP	Segment volume
ROI	Effectiveness

Table 2. TV curation as an advertising market, terminology equivalence. (Segment volume and effectiveness will be defined later)

Under this perspective, TV curation is then a supply-demand market where the supplied goods are curation slots and the demand role is played by content "aspiring" to be promoted. As we assumed that inventory (curation slots) is fixed, the supply curve is actually a vertical line.

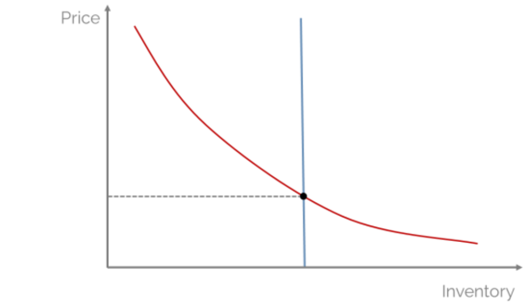

Fig 6. TV curation supply-demand curves.

Like with TV advertising, an analogous phenomenon of demand saturation can be observed: curating the same content in more than one position within the UX will typically yield no further results, so each segment is saturated for a particular content entry when this is curated once.

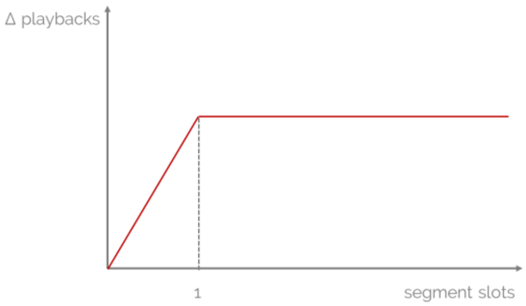

Fig 7. Hard-hit demand saturation in TV curation.

Inventory valuation

As part of our program for the interpretation of TV curation as an economic market akin to TV advertising, we need to find a numerical formula for inventory valuation that captures the utility of placing a particular content entry in a segment curation slot. This is a very straightforward proposal:

value(content→segment) = CM(cont) × volume(seg) × match(cont, seg) × WTR(cont,seg)

- **CM** is the "economic" contribution margin of each new user playback as a result of the content being promoted. We will analyse this quantity later on.
- The **segment volume** is defined as the number of unique user sessions in the segment for the curation period considered. This is akin to RP/TRP figures for audience reach calculation in TV advertising.
- **match(content, segment)** measures the degree to which the content is thematically aligned to the preferences of the target segment. For instance, the movie "Frozen" is a high match for a kids segment, but probably not so for an audience interested in, say, documentaries.
- **WTR** is short for **watch-through rate**, the percentage of impressions of the curated content that translate to an actual content playback by the user, much like the analogous **click-through rate** from online advertising. As thematic match has been factored out above, WTR correlates positively with the intrinsic "quality" of the content; also, it is expected to correlate somewhat *negatively* with the content entry popularity —if a content is very popular, users will look for it proactively, resulting in a portion of total playbacks not being funnelled through the curation slot.

To summarize, the value of placing a content entry in a given segment curation slot is simply an estimation of the expected incremental playbacks for the content, multiplied by their economic contribution margin.

Contribution margin

In microeconomics, the contribution margin of an individual product sale is just the revenue brought in by the sale minus the marginal cost of production. We can follow then this same definition in our scenario:

CM(content) = unit revenue - unit cost

From the standpoint of an editorial team in charge of curation, "revenue" and "cost" can mean actual monetary values (as in the simple case of pay-per-watch **transactional video on demand** or **TVoD**) or else be a proxy for some measure of service engagement, growth, etc. The following is just an approximation among many possible ones:

- Unit revenue:
    - Unit price for TVoD or similar pay-per-watch content.
    - Normalized unit price for subscription-based content (for instance, price of the subscription price divided by the average number of contents watched by a user).
- Unit cost
    - Unit cost for TVoD. If content acquisition operates under a **minimum guarantee** clause, this is zero until consumption reached the minimum.
    - Normalized unit cost or zero (sunk cost) for fixed-price content packages (for instance, a TV channel with an associated catch-up catalogue).

The fact that the unit cost can be zero (that is, we are operating under a sunk-cost assumption) allows us to model the typical situation where an editorial team wishes to promote content beyond its intrinsic qualities so as to reach the pre-agreed minimum number of playbacks.

Inventory pricing and allocation

In a market where the supply is scarce and fixed, a commonly used procedure for price determination is by **auctioning**. Furthermore, a basic result from Auction Theory says that we have an equilibrium strategy when all potential buyers (here, content entries competing for a curation slot) bid proportionally to their estimated value for the good being auctioned (with a fixed factor dependent on the number of buyers). With this in mind, we can simply state that the price each content is willing to pay for a particular segment slot is exactly **value(content→segment)**: this is also consistent with the editorial team's general goal of maximizing curation value.

Once we have derived a formula for inventory valuation and a procedure for inventory pricing, a very simple algorithm can be devised for **automatic curation**:

Max gradient algorithm for automatic TV curation

For each curation slot:

From contents *not yet curated* in the segment (demand saturation):

Assign content with maximum valuation

Adjust unit costs

This algorithm is not provably optimum but will be very close in practical situations.

Curation effectiveness

Much like in TV advertising a campaign is assessed for ROI, we can measure the **effectiveness** of a curation action by comparing its "economic" output against the value paid for the slot:

$$\text{effectiveness}(\text{content}\rightarrow\text{segment}) = \frac{\Delta\ \text{playbacks} \times \text{CM}(\text{content})}{\text{value}(\text{content}\rightarrow\text{segment})}$$

$$\Delta\ \text{playbacks} = \text{actual playbacks} - \text{baseline playbacks}$$

Plugging the definition of value(content→segment) in the equation, we get:

$$\text{effectiveness}(\text{content}\rightarrow\text{segment}) = \frac{\Delta\ \text{playbacks} \times \text{CM}(\text{content})}{\text{CM}(\text{content}) \times \text{volume}(\text{seg}) \times \text{match}(\text{cont,seg}) \times \text{WTR}(\text{cont,seg})} =$$
$$= \frac{\Delta\ \text{playbacks}}{\text{volume}(\text{seg}) \times \text{match}(\text{cont,seg}) \times \text{WTR}(\text{cont,seg})}$$

That is, the effectiveness is just the number of actual incremental playbacks divided by our priori estimation based on segment volume, content match and watch-through rate.

The quantity Δ playbacks depends on the determination of baseline playbacks, that is, the number of content playbacks that would have occurred nevertheless if the content had not been promoted. This can be estimated in a number of ways:

- Count as baseline playbacks those that happened outside the curation slot. This is just an approximation as the very presence of the curation slots "attracts" users who could have played the content through some other UX path.
- Use the number of playbacks for the content in the segment before curation happened.
- Use the number of playbacks for the content in **another** segment where the content entry was not promoted. This estimation is prone to be very rough as the number has to be normalized for the current segment:

  $$\text{baseline playbacks}(\text{seg}_1) = \text{playbacks}(\text{seg}_2) \times \frac{\text{volume}(\text{seg}_1) \times \text{match}(\text{cont,seg}_1) \times \text{WTR}(\text{cont,seg}_1)}{\text{volume}(\text{seg}_2) \times \text{match}(\text{cont,seg}_2) \times \text{WTR}(\text{cont,seg}_2)}$$

  which involves match factors and WTRs, estimated quantities on their turn with a high level of uncertainty.
- Do A/B testing on the curated segment.

Critical components of the model

The ability of the model to predict curation performance rests on two parameters measuring the degree in which content aligns with audience preferences, namely **match** and **WTR**. Doing a precise estimation of these, then, becomes essential in practical applications:

- Estimating the matching factor match(content, segment) is a problem similar to that of **personal recommendation** for the case of individual users, with the added complexity that recommendation engines typically resort to **collaborative filtering** to detect similar user profiles, a technique we cannot apply in our aggregated scenario —basically, collaborative filtering works with thousands of users and here we only have a handful of segments. So, matching must depend on an (informal or automated) analysis of the thematic qualities of the content as compared with the assumed preferences of the segment audience. **Content microgenres** become extremely useful here, as do techniques for automatic content categorization and clustering based on **automatic tagging**.

- As for the watch-through rate, we can further split it into:
  - A structural component that measures the average effectiveness of curation slots for a given segment, regardless of the content being promoted. This can be calculated very easily based on past history.
  - A non-structural component which is, by definition, highly correlated with the nature of the content itself. Estimating this component is probably best left to humans, who can bring in external knowledge about content popularity outside the service, social media buzz, etc., though automated supporting tools can be implemented for this task (**social trends analysis**).

Assisted TV content curation by Optiva Media

Background

Founded in 2002, **Optiva Media** is devoted to helping TV companies around the world make the most of their services. The company's services include all areas of the TV space from engineering to operations, media and metadata management, product design and development, research, and business and technical consultancy.

Optiva Media has developed **MediaSuite** (mediasuite.tv), an ecosystem of interoperable products for helping TV operators manage their content operations both at the metadata and media levels. As part of the evolution of **MediaSuite,** we are working on a new component called **MediaQurator** for the assistance of editorial teams with content curation. **MediaQurator** rests on the theoretical foundations explained in this paper.

MediaQurator

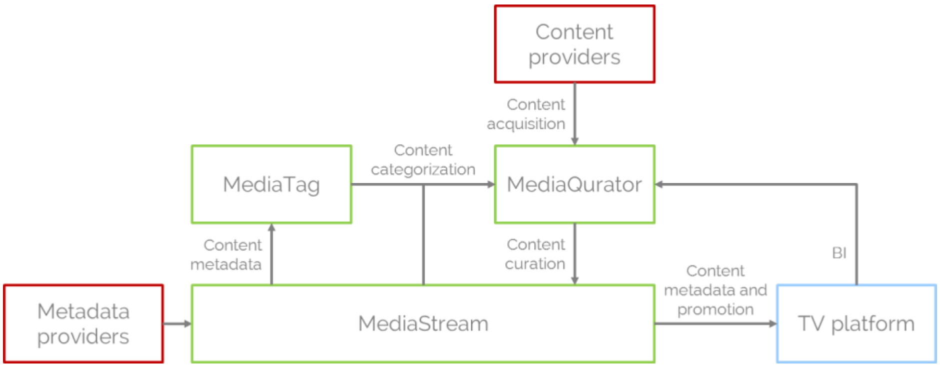

Fig 8. MediaQurator as a component of Optiva Media's MediaSuite.

MediaQurator develops the concept of Assisted TV Content Curation through the implementation of automated tools to help editorial teams to:

- **Plan** and **decide** on content promotion and placement.
- **Predict** the effectiveness of their decisions.
- **Measure and compare** actual results vs initial estimations.
- **Update and optimize** their curation strategies.
- Guide their **content acquisition** decisions beyond and prior to content curation itself.

MediaQurator interacts with the following **MediaSuite** components and external systems:

- **MediaStream**, an integrated metadata aggregation and handling system. **MediaStream** can provide metadata information on curated content to **MediaQurator** (mainly, genres) and accepts back lists of curated content as decided by the editorial team.
- **MediaTag** automatically extracts thematic tags from content metadata using AI algorithms. As discussed earlier, this is valuable for assessing content match against a given segment audience.
- **MediaQurator** integrates with the BI subsystem of the customer's TV platform to obtain information on predefined segments and service metrics used to gauge curation effectiveness.

Further investigation

Starting from this foundational analysis, we would like to investigate on a number of open tracks:

- The calculation of playback contribution margins for subscription-based content can be refined to better understand the economic gain attributable to an increase in usage for these types of services.
- The current approach to inventory valuation has centred on the actual monetary value directly or indirectly brought about by content playback. Alternatives exist that focus on other KPIs of interest to the editorial team such as user engagement, churn, etc. Hybrid formulations can also be devised.
- There is an entire subdomain of investigation related to calculation of content matching through automatic tagging, content clustering, analysis of social media trends, etc.
- The model makes the simplification of assuming that each slot (within a given segment) is basically interchangeable. There are curation scenarios, however, where thematically related content is grouped in **curated playlists**. It remains to be studied how this curation strategy increases engagement with respect to the baseline scenario.
- Along the same line, the notion of **DJ playlists**, curated content lists whose authors are publicly known and have a following of their own, introduces a new engagement factor with its particular economic dynamics.
- If we look upwards the value chain of a TV service, editorial teams are also typically responsible for **content acquisition**, which can be amenable to theoretical study. Market models can be devised for this activity where the roles of supply and demand are swapped with respect to the TV curation model we have just presented. Furthermore, combining both acquisition and curation can provide an even more comprehensive view of the role of content aggregators in the modern-day TV industry.