



AUDIENCE 2016 MEASUREMENT

**Is this the Beginning of the End
for Age and Gender Targets?**

ARF Whitepaper

Pete Doe, Chief Research Officer

1. Introduction

Why is TV still bought using age and gender targets? This is a question that has been asked for many years, and there have been many calls for more specific audience based targeting¹. In the last five years, data sources have proliferated and this has enabled more detailed understanding of consumers' viewing patterns (and equivalently, viewers' product consumption and attitudes). So in a world where big data and digital addressability are standard why do age and gender targets persist and is this beginning to change?

This paper attempts to outline how linear TV is beginning to exploit data in the buying and selling of advertising. It covers the following topics:

- The opportunity – using real-world examples, we show that there are significant gains delivered by more precise data-driven targeting of linear ads
- How technology is employed to create schedules with advanced targets
- The types of data sets that are available in the US for advanced audience targeting and their pros and cons
- Methodological challenges that need to be addressed when using advanced data sets in linear TV deals.

So what do we mean by advanced audiences or advanced targets? A simple answer is anything that goes beyond traditional age and gender targets. This can range from more detailed demographic or market break information such as income or presence of children to brand specific customer data for a specific advertiser.

It's worth noting as well that this paper does not use the term "programmatic TV". This term is a source of confusion and misunderstanding for many, and has connotations of remnant inventory, black box "tech tax" and addressability via real time bidding systems. These connotations are mostly irrelevant in the context of this paper because linear TV is very different from digital. Key differences are:

- Linear TV is finite in terms of inventory while digital can be viewed as effectively infinite
- Linear TV still has a mass reach capability for brand building that is much more difficult to achieve with digital
- Addressability is still a small part of linear TV, though it will increase in the coming years as the technology becomes more prevalent.

What advanced audience targeting offers in linear TV is better use of an advertiser's budget, better valuation of media owners' content and more relevant ads for viewers.

2. The Opportunity

The examples in this section are based on actual advertising schedule proposals created for clypd clients.

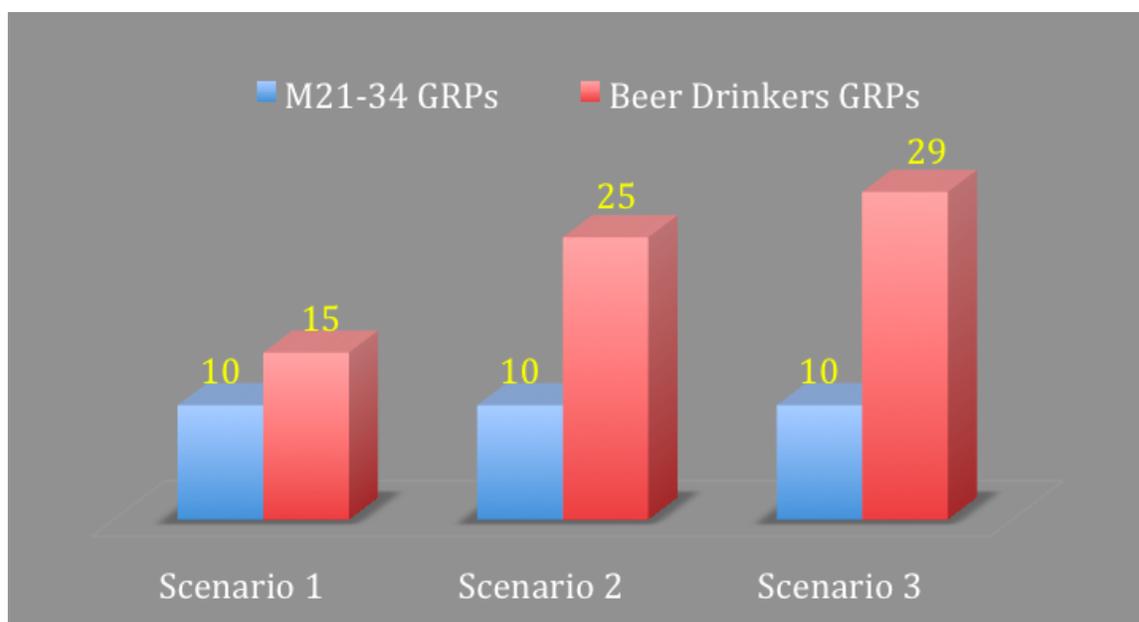
Example 1. A beer advertiser wanted to improve its effectiveness, defined as improving its effective CPM for beer drinkers by changing its normal mix of networks and day-parts in a deal with a media owner, for a Q3 2015 deal.

In this case three scenarios were created, all designed to deliver the same demographic impressions (10 GRPs for Men 21-34) with the same budget. Scenario 1 was the advertiser’s typical deal, with just two out of six networks being selected and most of the GRPs in prime.

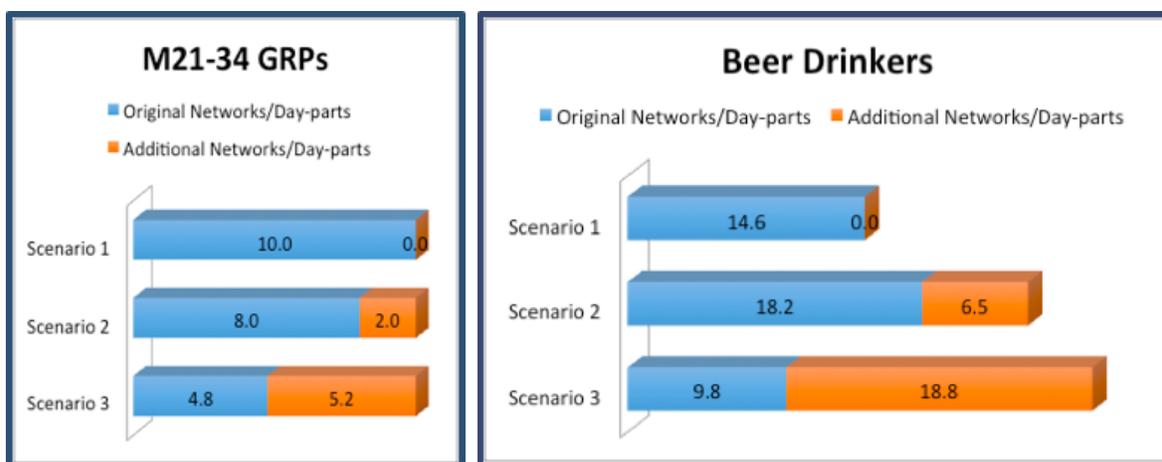
Scenario 2 relaxed the network constraints but kept the same day-part mix while Scenario 3 relaxed both network and day-parts. In all three cases the schedules were created to optimize the number of advanced target impressions, with the advanced target in this case being “heavy beer drinkers”, defined using the Nielsen GfKMRI data fusion.

	M21-34 GRPs	Optimize Advanced Target?	Networks (Dollar Mix)						Day-parts (Impression Mix)			
			A	B	C	D	E	F	Fringe	Week-end	Prime/Late Night	Over-night
Scenario 1	10.0	No	95%	5%					15%	15%	70%	
Scenario 2	10.0	Yes	75%			25%			15%	15%	70%	
Scenario 3	10.0	Yes	50%			50%			10%	20%	20%	10%

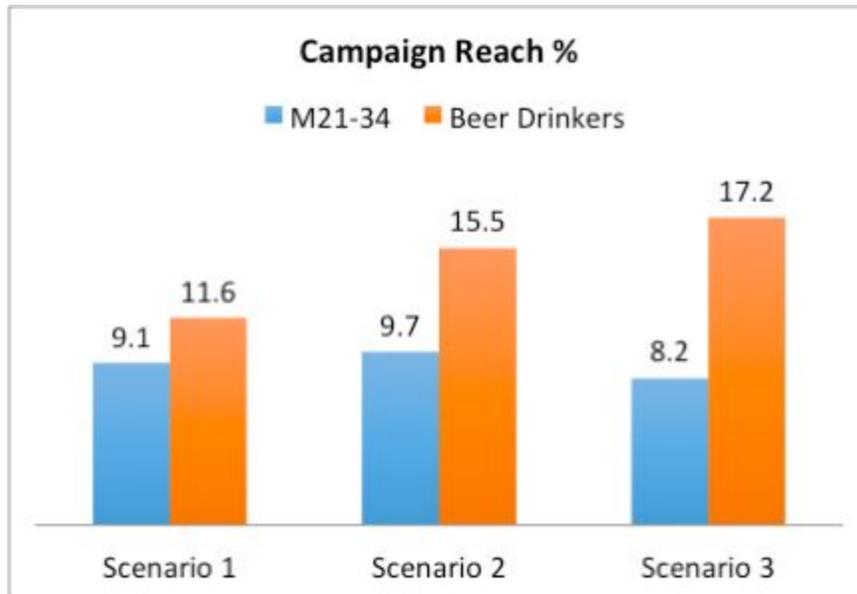
So how did they perform? Scenarios 2 and 3 both delivered significantly more beer drinker impressions. Scenario 3 performed better than Scenario 2, demonstrating that the more relaxed the deal constraints, the better the delivered audience. In the case of Scenario 3, nearly the beer drinker impressions were nearly doubled, or equivalently, the effective CPM for beer drinkers is nearly halved.



When we examine the scenarios in more detail it is interesting to see how the GRPs differ from the original networks and day-parts in Scenario 1. Scenario 3 has less than half its M21-34 GRPs coming from the original networks and day-parts. For beer drinkers, the increase in GRPs comes from both the original and additional networks and day-parts: the inventory selected in Scenarios 2 and 3 deliver a higher concentration of beer drinkers than Scenario 1, and the remaining budget is shifted to additional network and day-part combinations that are popular with beer drinkers.



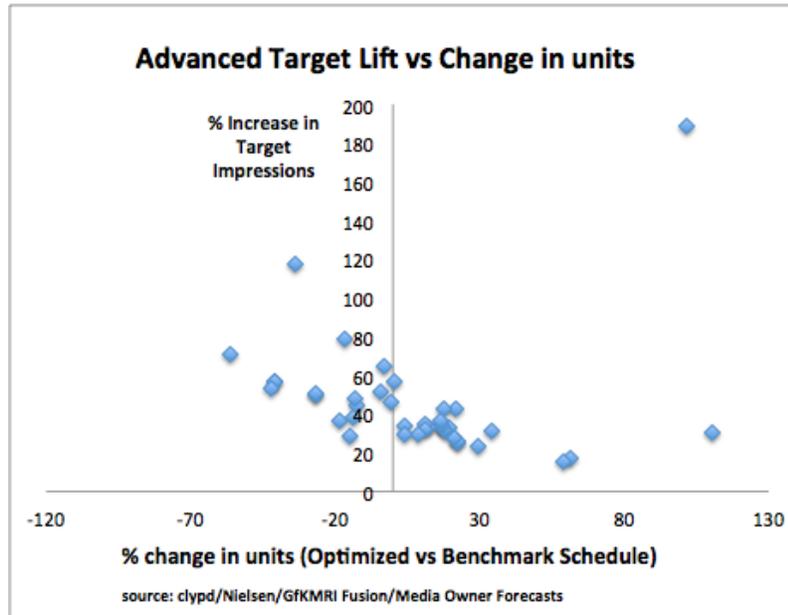
In delivering more beer drinker impressions, a question also arises around reach and frequency: to what extent are the primary demo and advanced target reach affected by the revised schedules? In this example Network A is the media owner's flagship network, and it delivers broader reach than the other more niche networks. Shifting away from Network A would be expected therefore to reduce reach. In the case of Scenario 3 this is indeed the case – M21-34 reach reduces from 9.1% in Scenario to 8.2%. However, for beer drinkers reach is increased in both Scenarios 2 and 3 – so the revised schedules are delivering both more reach and frequency for the advanced target. An upcoming development of the optimization algorithm will include reach constraints as an option, and that may be a neater way of developing schedules with minimal constraints on networks and day-parts/selling titles but reach constraints in place to prevent a completely frequency-driven solution being created.



Another key dimension of the deal is number of units. In this case the Scenario 2 and 3 deals use more units than the benchmark Scenario 1 deal. This is to be expected as the schedules move away from the higher rated network A and (in Scenario 3) primetime slots, while maintaining the budget and demo impressions delivery. These deals have therefore identified less expensive inventory that nonetheless delivers valuable advanced target impressions. Knowing this, the media owner can better value inventory.

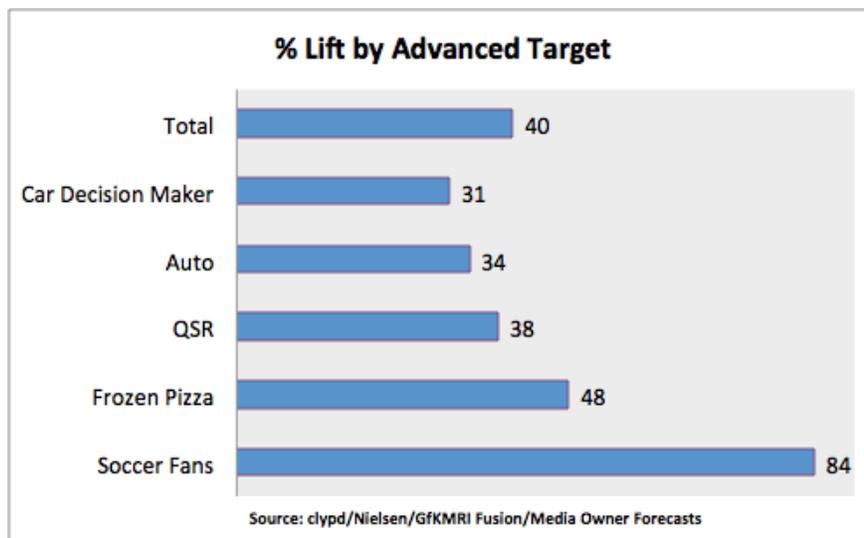
Number of Units in Deal	Networks					Total
	A	B	C	D	E	
Scenario 1	16	84				100
Fringe M-F 4p-6p	5	1				6
Prime M-Su 6p-12a	5	2				7
Late M-Su 12a-3a	1	81				82
Weekend Sa-Su 9a-6p	5					5
Scenario 2	28	42	6	68	6	150
Fringe M-F 4p-6p		6				6
Prime M-Su 6p-12a	18	20	4	38	6	86
Late M-Su 12a-3a	6		2			8
Weekend Sa-Su 9a-6p	4	16		30		50
Scenario 3	21	30	71	151	6	279
Fringe M-F 4p-6p		6		2		8
Prime M-Su 6p-12a	11	6	44	80	6	147
Late M-Su 12a-3a	6		14			20
Weekend Sa-Su 9a-6p	4	16		56		76
Overnight 3a-6a		2	13	13		28

In other cases we have seen the number of units reduce as the advanced target is found disproportionately in higher rated and more expensive shows. The scatter plot below shows advanced target lift ranging from about 20% to 180% with the number of units varying from 60% fewer to over 100% more.



This scatter plot uses results from over 80 campaigns, across categories including Pizza Lovers, QSR visitors, Soccer Fans, Car Decision Makers and Auto Insurance Switchers. In all cases the deals were constrained by budget and demographic CPM for benchmark and optimized campaigns

In summary, the optimization produces significant lifts across a variety of categories. In this set of targets the average lift was 40%.



3. How it Works

Advanced target deals are being developed in various ways: one common way is described below. With this approach it can be seen that there are traditional discussion and review steps between buyer and seller – this is not a completely black-box experience by any means. There are other scenarios that could be considered – for example, buyers may wish to create their own schedules and send these as orders to the seller. In that scenario, the media owner would need to open up available inventory and current rates to the buyer, potentially losing control of the process. Game theory suggests that the best outcome in these circumstances is complete transparency on both sides but commercial reality is not an academic exercise: that said, the best deals are usually made in a spirit of co-operation and open-ness where both sides feel that they have won. With advanced targeting there should be three winners – the advertiser/agency who have achieved their goals more efficiently (ie with a lower cpm), the seller who has better valued inventory, increasing yield, and finally the viewers (sine qua non) who should be watching ads that are more relevant to them.

Step 1: Agency/Media Owner agree advanced target

This step is sometimes lengthy as there are many ways of defining an advanced target, with different data sources and classifications within those data sources. This is perhaps one of the reasons that age/gender targets have persisted for so long: a currency measurement has both a value element and a label (eg \$50 has 50 and dollar) and for TV that has long been defined as a number of demographic GRPs (eg 100 Men 18-34 GRPs). The simplicity of this has been its strength – age/gender targets are clearly understood and available in most research data. With advanced targets there is less clarity. Conceptually a target may be clear – the target is international travelers for example – but what do we mean by an international traveler? Traveled in the last year or ever? How many times? Does Hawaii count? What about a day-trip to Canada for someone who lives near the border? What data should be used to define the target and how can it be activated against TV inventory? Having some common framework to define advanced targets is a crucial step in the process.

Step 2: Deal parameters are defined

This includes dates, budget, cpm/impressions/GRP goals and network/day-part/selling title mixes. It also likely includes a demo guarantee as well.

Step 3: Deal proposal is created

The schedules presented in the previous section were created using a technology platform installed in the media owners' sales systems. This technology comprises:

- a) A user interface for defining requirements of the deal(s). This can be private within the media owner or exposed to approved buyers either with manual input of requirements or via an API.

- b) Data access to audience rating data (including advanced target data) and media owner inventory and rates. These audience rating data include media owner and/or clypd-created forecasts of impressions: usually the media owners have forecasts for age/gender targets but not for advanced targets, so these are created automatically by clypd within the platform.
- c) Algorithms that combine the deal requirements, audience data (including advanced target audiences), media owner inventory and rates to create deals to buyer/seller agreed specifications. These specifications typically include the following:
 - The budget
 - The campaign/flight period
 - The available inventory and rates
 - CPM/Impressions/GRP goal for a demographic target
 - Deal parameters of networks, day-parts/selling titles and weekly weights (this includes inclusion and exclusion rules across these dimensions)
 - Separation constraints (eg no more than two units per hour)
 - Spot length
 - Crucially, an advanced target that is to be optimized for delivery within the above constraints

In addition, some optimizations require the re-casting of an existing deal with the constraint of number of spots, though this is usually unconstrained.

- d) Connection to media owner booking system for activation of the deal. At this stage there may be some additional tweaking of the schedule to conform to requirements around separation from other advertisers or creative conflicts.

Note that these deals typically still include a demographic (age/gender) guarantee. Buyers and sellers are exploring deals based solely on advanced targets but as most deals have been agreed up-front based on age/gender, optimizing an advanced target within the pre-agreed demo deal constraints has been seen as a good way to proceed.

Creating an advertising schedule that satisfies all the requirements above is a difficult and time-consuming task for someone working with basic spreadsheet capabilities but it is a quick solve for the algorithm. To assess the value of the proposed deal the algorithm also generates a benchmark schedule that satisfies all the deal requirements but without optimizing for the advanced target. This gives a very clear demonstration of the value of the optimized deal in delivering more advanced target impressions. There are also other benchmarks that buyers and sellers may wish to use such as last year's deal or a more general run of schedule comparison, ignoring inventory availability.

To date, deals and proposals have been developed on a case by case basis but the algorithm includes functionality to cover multiple deals and even entire media owner inventory pools if required.

Step 4: The proposal is reviewed by the Agency

This results in acceptance or some changes/tweaks leading to a final deal being agreed.

Step 5: The deal is activated

At this stage the deal is live and runs through the same systems as a traditional deal. However, for stewarding and posting of advanced targets additional reporting capabilities have to be created. The issues that need to be addressed here are discussed in section 4.

4. Data Sets

Data about consumers comes from many sources these days, including traditional consumer surveys, data and marketing clouds, proprietary data from media owners, agencies or advertisers. For a data set to be viable for advanced targeting of linear TV the following is necessary:

Permissible Use: the data must be legally available for use in the deal. All ownership, licensing and privacy issues need to be resolved.

Relevance: The data must be able to adequately define the target of interest. Often a perfect description of the target is unavailable so the key consideration should be “does this data define my target better than a standard age/gender target?”

Connection to Currency: the data must be connected to linear TV inventory to enable the value of the inventory to be assessed with sufficient credibility for buyer and seller to be able to transact. This means that inventory units have a price and (usually) an audience estimate. The agreement between buyer and seller concerning what data to use in the transaction is just that: a decision between buyer and seller, and for these deals to be conducted, systems suppliers like clypd need to be data-agnostic, supporting whatever data sets are required. For most linear TV deals in the US, Nielsen National People Meter data are the source of the audience estimates and the advanced target data are often connected to this source. However, it is feasible to create a deal using Nielsen for a demo target and another data source for advanced target.

Persistence and Timeliness: The data must be available in a timely fashion and persist long enough to enable planning, activation, posting, and if required, effectiveness measures such as ROI estimation. A single set of data giving a snapshot of consumer target performance from three years ago is unlikely to be usable.

There are many types of data sets that satisfy these criteria, and once agreement is reached between buyer and seller there are some technical details around the data that need to be defined and addressed clearly and consistently. These include:

- a) **Data Access:** Ingesting the data in a stable and timely way and enabling its use in activation is a non-trivial task. Systems suppliers end to be nimble and flexible as timelines for deals are often short.
- b) **Demo Qualifiers:** Of course, the point of advanced targets is to go beyond demographics but there may be reasons to include a demographic qualifier, eg Adults 21+ beer drinkers in the US, rather than Adults 18+ beer drinkers.
- c) **Classification Gaps:** If the data are connected to metered panel data such as Nielsen's NPM, what calculation procedures are needed to address classification lag or other non-classification issues? The daily cadence of reporting and panel updating does not typically apply to fused or matched data sets, meaning that gaps in the data occur.
- d) **Consistent Universe Estimates, GRP and Impression Calculation:** Audience measurement systems are typically geared towards demographic reporting that ensures a clear and consistent view of universe estimates. If you sum the weights of the Nielsen panelists aged 18-49 you will get the same results for every day of a broadcast year – the Adults 18-49 Universe Estimate. This does not happen for advanced target data as panel systems do not and cannot control thousands of consumer targets at that level. Calculation procedures therefore need to accommodate this variability in a consistent and transparent way.
- e) **Visitor Viewing:** In people meter systems, visitors in panel homes are included in the measurement as a surrogate for panel members' viewing in other homes, and visitors are classified by gender and age. Audience estimates for gender/age groups such as C3 ratings include the contribution of visitors to the audience. In contrast, it is usually impossible to classify visitors in terms of advanced targets since there is insufficient information about the visitors to support the matching process. For this reason, analyses of advanced targets typically exclude visitor viewing and this can significantly under-estimate the advanced target audience – visitor viewing is about 5% overall but can be as much as 20%, particularly for younger skewing audiences. To address this issue, it is possible to estimate the contribution of visitors to advanced target audiences using a scaling/factoring approach. It is therefore important to define at the beginning of any deal whether the advanced target will include visitor viewing.
- f) **Reach and Frequency Calculation:** For respondent level classifications, advanced target reach and frequency calculation is fairly straightforward, but consideration should be given to classification lags and visitor viewing: probability models can be employed to address these issues and ensure an unbiased reach estimate is obtained. For advanced target data sets that are based on aggregated data feeds
- g)

reach and frequency cannot be calculated directly and some strategy to model reach needs to be agreed.

Audience Forecasts and Consistency: Predicting audiences is a challenge, whether for demographics or advanced targets. When creating advanced target forecasts the

predictability and consistency of the data will inevitably be less than for a broad demographic target. And as linear TV is a great way of building brand awareness, “hyper-targeting” a very narrow segment is not a good strategy anyway. Equally, creating a very broad target reduces the value of the target for the advertiser – a sweet spot needs to be found. Another aspect that needs to be considered is that high indexes in one period tend to reduce in future periods. A program that indexes at 200 for an advanced target in a planning period may only index at 150 in the future. This tendency is a naturally occurring phenomenon known as regression to the mean (it is one explanation of the “Sports Illustrated Cover Jinx”²) and what it means for campaign schedules is that inventory selected based on high historical indexes will likely perform less well than the initial forecast suggests in delivering a high index. This can be mitigated with more conservative forecasts and in-flight stewarding, and should definitely be considered when managing expectations around a deal and any guarantees.

Data Sets Overview: Some data sets that are being commonly considered and used when creating deals are described below. Each has pros and cons. A common feature of all of these data sets is that none of them was created for linear TV currency transactions. They were built for planning, ad sales or effectiveness and the requirements of these use cases are less stringent than currency requirements. As a result, there are issues that need to be considered and addressed when using these data in actual deals, particularly if sales guarantees are agreed based on the data.

NPM/MRI Data Fusion

The NPM/MRI fusion overlays profile points from GfKMRI’s “MRI” study onto Nielsen’s National People Meter TV viewing panel (“NPM”) to enable planning, ad sales and advanced target advertising across a wide range of consumer attributes.

This data fusion overlays MRI fields onto NPM using a statistical matching process. Each NPM respondent aged 18+ is classified with 20,000+ MRI profile points in these categories:

Auto Transportation	Financial	Magazines
Baby Needs	Food	Media
Beverages	HBA	Pet Needs
Biz Purchase	HH Items Non-food	Psychographics
Candy	Home Improvement	Shopping
Clothing Accessories	Influentials	Tobacco
Consumer Electronics	Insurance	Toys
Dr Visits	Job Title	Travel
	Leisure Activities	
	Restaurants	

This data fusion is a statistical integration based on respondents' common characteristics, not name and address matching. Because both data sets are samples, the overlap between the two (ie people on the NPM who also took part in the MRI survey) is minimal, preventing direct matching.

Pros of this data set are:

- The data sets used in the fusion are currency quality and used throughout the industry by both agencies and media owners. The fused data set is also well-used in the industry. This provides a great starting point for agreeing an advanced target – there is common ground and understanding between buyer and seller.
- The fusion ensures that all persons 18+ are included in the matching. This is not possible with name and address matching of data sets, where match rates typically range between 50%-80%
- The data are representative of the population, using high quality sample designs.
- The respondent-level classification is very flexible for analytic purposes and enables reach and frequency analysis.

Cons/Concerns that may exist with this data set are:

- The MRI survey is based on claimed past behavior and attitudes, not actual recorded purchase behavior.
- Model bias: the data fusion uses a lookalike modeling process that inevitably contains the potential for bias.
- Classification lag: the fusion classification occurs monthly in arrears and the Nielsen panel is updated daily. At times, up to 20% of Nielsen panelists may be unclassified with MRI fields. That issue needs to be addressed to enable in-flight reporting and stewarding, and timely posting of campaigns.
- ROI measurement that directly measures ad exposure and purchase, pre-and post campaign, is not enabled by these data. While technically possible, the data set does not have sufficient statistical validity at this level.

Nielsen Buyer Insights (NBI)

This data set matches credit and debit card transaction data to the Nielsen panel via a third party name and address match. The credit card data are aggregated across banks and cover about 80% of credit card activity and 20% of debit card activity.

Pros:

- The data are real transaction data, not recalled behavior.
- The direct matching enables ROI analysis, allowing assessment of ad effectiveness

- Credit cards are used widely in the US across a wide variety of categories eg airlines, department stores, financial services, movies, wireless.
- Respondent level data are available for flexible analysis and reach and frequency.

Cons:

- Coverage is incomplete: some NPM panelists are unclassified, either because they do not use the credit cards in the data set or because the name and address data are unable to be matched by the third party matching agency. Because some groups (eg Spanish speakers) are more likely to use cash than credit cards, this coverage skew can result in a skewed media plan.
- The detail from credit card records is at the retailer level. A transaction record related to a department store is the same whether the purchase id for apparel, bedding or cosmetics.
- There can be some inconsistency between the person using the card and the name on the card.
- As with MRI fusion, there is some classification lag that needs to be addressed if the data are being used in-flight and for timely posting.

Nielsen Catalina Solutions (NCS)

This data set matches frequent shopper data from co-operating retailers with Nielsen panel data and set top box data.

Pros:

- The data gives brand and category level detail for CPG
- Direct matching enables ROI measurement
- Incomplete matching is addressed via calibration of viewing to Nielsen currency data, though this applies to aggregated data feeds, not respondent-level.

Cons:

- The coverage is constrained by retailer co-operation agreements.
- The data relates only to the CPG category.
- As with MRI fusion and NBI, timeliness of data update

Other Data Sources

More and more data sources are available for better decisions to be made in linear TV. Comscore (formerly Rentrak) and TiVo Research Data provide large data sets

based on return path data that are linked to a variety of consumer data in various categories, including industry standard vertical data such as Polk auto purchase data, consumer segmentation systems such as Acxiom Personix and linked CRM data sets. The potential here is obvious, provided buyer and seller agree to the use of the data (See Requirement 3 above). To date the acceptance of these data sources as viewing currency has been relatively limited, since the undoubted benefit of large samples compared with Nielsen has not outweighed concerns about representativeness and the measurement of homes/devices rather than individuals.

In addition to these, agencies and advertisers have their own data sources, which can be linked to viewing data through statistical integration and/or data matching. In some cases the buyers may wish to keep the details of these targets secret from the media owner and there is theoretically no reason why the seller needs to know everything about the target. However, some transparency is essential if deals are being guaranteed, to ensure that the deal is stewarded and posted fairly.

Media Owners are also creating data capabilities and creating consumer segments. It will be interesting to see whether this gains acceptance for broader deals – would Media Owner B accept a deal based on an Advanced Target created by Media Owner A? For this to happen the advanced target data source would presumably need to be scrutinized very closely by an independent expert to make Media Owner B comfortable that their inventory was being fairly valued.

5. Conclusion

This paper has focused on the promise of advanced targets in linear TV, in particular the data sets that can make this promise a reality and the challenges that advanced target data bring. There are many other aspects to this that are outside the scope of this paper – technology standards that ensure buyer and seller systems can connect, the building of system to ensure data flow from data suppliers through deal creation and into booking systems, system integrity checks on all the above, and external verification that the value of advanced targeting is being realized. Integrating linear TV deals with digital is also another topic that demands a whole paper.

But the importance of data cannot be overstated. Age and gender targets have persisted in currency transactions for all these years because they are simple and available: they provide a common and unambiguous point of agreement between buyer and seller. Advanced targets inevitably increase the complexity of the discussion, but offer gains for buyers, sellers and viewers. For the full potential to be realized there needs to be collaboration across the industry, particularly between media owners who stand to gain most by presenting unified capabilities to advertisers who understand and value linear TV.

Notes

¹For example: <http://www.cbsnews.com/news/are-you-targeting-the-wrong-audience/>

² https://en.wikipedia.org/wiki/Sports_Illustrated_cover_jinx