A Guideline from TAM for all TAM Data users

I. Introduction

TAM data fulfils the need to understand Indian television audience behavior. Over the past decade, the data and its dimensions itself have vastly increased; more sample size, more geographies to analyze (28 reported markets in 2008 to 42 reported markets in 2013), platforms like Digital emerging, explosion in number of active channels 427 in 2009 to 509 in 2013 in etc.

The increase in data has also led to increased data analysis intensity. Unfortunately, there are cases – and increasing in frequency – of data being dredged to yield interpretations that are not meaningful but merely force-fit into a pre-determined conclusion. Some use creative means to represent data which only results in misrepresenting or misstating facts. All of this puts pressure on the rest of the industry players as well as the audience measurement agency.

The Indian TV environment is going through unprecedented change. Witness the transition to digital platforms by millions of analog households owing to DAS. Changes like these in the environment lead to changes in viewing behavior which is in turn reflected in the data. In this period of change, it becomes even more important to be careful in interpreting and representing data.

This note briefly prescribes guidelines that a TAM data user should keep in mind before analyzing and representing data or seeing the products of such efforts. It does not claim to be exhaustive but covers the most common areas on data (mis)representation. Some ideas in the note are presented as if written with a data analyst/representer in mind and others with a viewer of the analysis/representation in mind. The context will make this clear but ideas are applicable to both sides of the representer-viewer equation.

The document has been divided into two main sections – one focusing on concepts and general guidelines and the other on graphical representation. This document is not about the best ways to represent data which is a separate and vast topic in itself. All data used here is purely illustrative and not based on actual TAM data. We hope that this note is taken to be a gentle well-intentioned guide to data interpretation and representation. It is not intended against any specific user or any groups. One of the dangers about writing a note like this is

that it itself becomes a repository of bad ideas to follow - we hope that does not become the case!

II. General Concepts and Guidelines

1. Realize that this data comes from a survey based on a statistical sample

Let's say you were analyzing a programme's performance for Mumbai for the C&S 4+ target group on June 11, 2012 that ran between 20:00 – 20:30 hours on Channel ABC. You ran the analysis in Media XPress and found that the reach of the programme was 5% and the associated sample size for the analysis target group was 2000 individuals. These 2000 individuals come from a set of homes chosen to be representative of Mumbai's C&S population aged four years and above who dwell in private homes with their family.

What if we had another set of 2000 individuals that we could survey at exactly the same time with exactly the same design using exactly the same process? Would exactly 5% of these individuals also watch programme ABC. Unlikely. But it is very likely that the number be close to 5%. Now take this further and imagine if we had a way to take millions and millions of such samples of 2000 from the population of Mumbai. Not all of these estimates will be close to 5%. In fact some may be off by a wide margin. We will never know the true proportion in the population but thanks to the statistical science, we do know that an overwhelming majority of the estimates (two-thirds) coming from these several million samples of size 2000 individuals will be between 4.3% and 5.7%. This range (4.3%, 5.7%) came to us from Appendix 1 which lists the standard errors for different sample sizes and estimate levels.

The moral is that any sample-based estimate is subject to the laws of statistics – and variation is one of them. One big implication is that very small differences in viewing between channels or programmes should not be taken very seriously – as the below point explains.

2. Understanding zero ratings - did no one really watch us?

This is a corollary to point 1 above. A question that comes up at times is of the following type: "Are you telling me that there is absolutely no one among SEC A 4-14 Males watched my show".

If the show was telecast at, say, 3 am then the answer can be "most probably, yes ". Note that we use the phrase "most probably" since we are dealing with a survey sample which has viewing probabilities associated with it rather than definitive ***Y***es-***N***o statements.

On the other hand, if the show was a show telecast at, say, 8 pm then we would say "there is always going to be some probability that your show delivered a zero rating even though there could well be a section of the population viewing your show". How do we make sense of this statement? Once again, we note that the rating is an estimate produced from a sample. For a sample of size n individuals drawn from a population of size N individuals, there is a non-zero probability that all n individuals did not view a show while some section of the N-n individuals did. This can happen only due to two reasons:

a. The actual viewership in the population is itself very small
b. n is small

For an idea, if the actual population TVR of a show was 0.5% and the analyst was analyzing data on a sample size of 50 individuals, there is ~80% chance that she will get a zero TVR in the output. But at a sample size of 1000 individuals, this chance of a zero rating dramatically reduces to just ~1%. Even a sample size of 200 individuals reduces the chance of a zero rating output (when the true population TVR is 0.5%) to ~37%.

Now assume that the analyst was looking at some estimates for a show whose true population TVR was 1%. The chances of her getting a zero TVR for this show at a sample size of 50, 200 and 1000 individuals is ~61%, ~13% and 0.004% - a large improvement over the above 0.5% case.

All this illustrates the importance of choosing reasonable sample sizes for analysis and also treating relatively small TVR estimates with care while analyzing them.

3. <u>Business significance is as important as statistical significance.</u>

Every standard statistics textbook will talk about the difference in estimates being 'statistically significant' and "practically significant".

Consider the graph below. At first glance it seems that Channel B is way ahead of Channel C and definitely ahead of Channel A. But a little attention to the

estimates themselves shows that the difference in estimates is very small: 0.01, 0.03, and 0.02 between Channels B-A, B-C and A-C.

Assuming these differences were even statistically valid, from a common-sense perspective ("practically significant") these differences are extremely unlikely to be large enough to impact business.

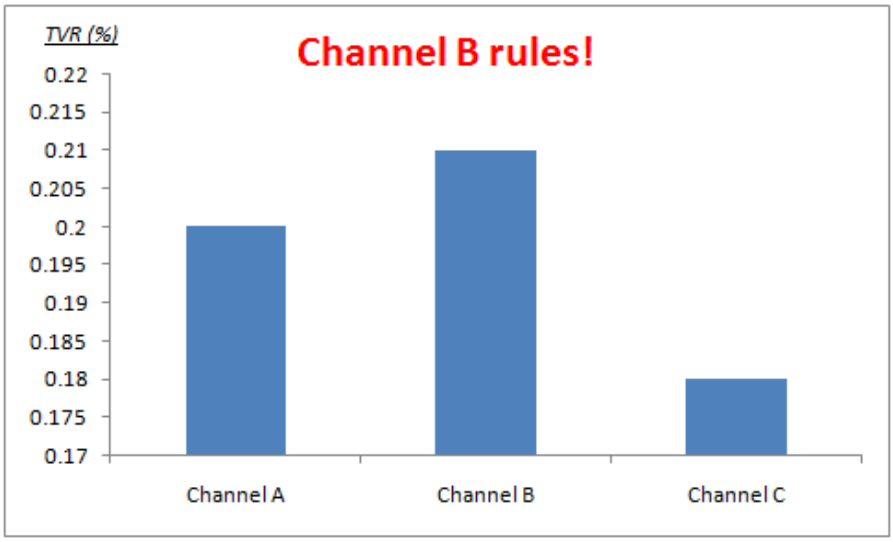

4. <u>Trend the data</u>

The biggest advantage of the TAM panel is that one is measuring more or less the same people (20%-25% average yearly churn) over time. Due to this, changes in viewing behavior are reported more precisely than if we were to simply conduct two separate surveys at different points in time. But the power of the panel is not about just calculating a change over some two points in time. We can see how gradually or suddenly this change is happening by also trending the data between the two points of time. And importantly, to judge whether this change was a 'one-off' or part of a systematic change.

Take for example a claim that states (only) that a channel's reach "has increased by 30%". Closer examination of the claim showed that two specific weeks were used to show this growth (Week 9 and 18). The claim as such is numerically correct. But the conclusion gets reversed completely when we look at the trend of the channel's performance and find that reach is actually in a declining phase as seen below:

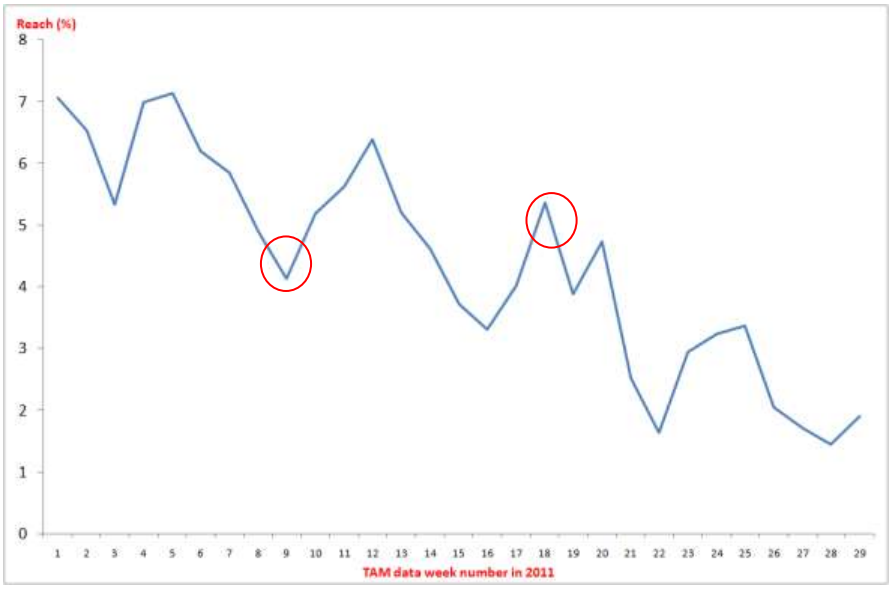

The lesson for a data analyst is pay close attention to trends and not on point estimates alone.

5. Use Averages

The above example also illustrates that it is not a good idea to look at any one given week/day/hour etc. in isolation to judge the viewership of a TV property. One must take averages to 'iron out' special changes in behavior to get a more 'representative' idea. For example, a box-office hit can certainly bolster viewership for a day but this special performance may be unlikely to be a regular feature of that channel.

Standard errors, which we referred to in Rule 1, get reduced when we take averages. By aggregating minutes into a reasonable time-band one is, in effect, conducting averaging and thus reducing statistical error. The minimum time element that can be analysed is a minute. But such an analysis will be associated with potentially larger statistical error. By isolating individual minutes or very few of them we are just increasing the scope for misreading data. Also, doing such analysis is very unlikely to be justified from a substantive business viewpoint.

How many data points should one take for averaging? This depends on the analysis context but as a guideline take at least 8 weeks if you are looking to conduct weekly averages.

6. Rankings can be misleading

It is fairly easy to represent oneself as a "number one" channel by isolating very specific dayparts/TGs/weeks where one got a relatively higher 'viewership'. But the real test is if the programme is able to reach out to a relevant TG in an effective, efficient and consistent manner. This varies by circumstance. It is quite possible that a programme ranked number 10 (measured, say on average weekly TVRs) attracts an audience profile that the 'better' ranked programmes have been unable to.

The bottom-line is that a complex phenomenon like TV viewership cannot be reduced to a simple ranking system. TAM, for one, never publishes ranking of channels or programmes.

7. State the estimate clearly

In a graphical context, this translates into labeling the axes clearly. It is not enough for a vertical-axis that is titled "audience size". Is "audience size" in the graph TVR in thousands? Or Reach is 000s? Or "Reach at 10 min+ continuous viewing"?

Another case is the often-misused weekly GRP estimate. By practice, many compute this as a half-hour GRP estimate but one could artificially inflate this estimate simply by choosing a smaller time interval for GRP computation. In any case, one should clearly state that the estimate is a "30-minute weekly GRP estimate". Do not assume that the viewer of the presentation, mailer etc. will know the basis of the representation. Point 8 below has more detail that one must provide in any communication.

8. State the estimate basis clearly

Apart from the estimate and its type, one should also state the estimate basis clearly. A data is a product of a set of analysis parameters. These should be stated clearly – in the header/footer, for example. This should include details of:

- Geographic markets chosen for analysis
- Time element used (Year, Month, Dates or the TAM data week)
- Target Group analysed
- Nomenclature used for e.g. does "kids" in your analysis stand for individuals in the 4-9 age group or the broader 4-14 age group? Does Late Night refer to the timeband starting 10:30 pm onwards or 11:00 pm onwards?

- Special filters. Are the reach estimates displayed in your chart based on a "5 minutes+" time spent filter?

  Please use humanly readable font for the header/footer. Conversely if you are a data user please pay close attention to the header/footer.

9. Watch out for 'selection bias' - Eye-grabbing headlines should not mislead

"Eye-grabbing" in the above phrase is not used pejoratively. In fact, our attempt in representing data should be to attract our audience's attention to some consumer insight. But we often find tables and charts that are summarized by an eye-grabbing title that does not match the representation itself. A simple example is a chart that is titled "we are the number 1 news channel in India" but based on a survey of only select cities.

10. Use indices judiciously

Indices are a popular form of data representation – often for the wrong reasons. The three questions we must ask yourself if you are using indices to represent data:

a. Have we boldly declared the actual numeric estimate used to compute the index?
b. Have we used a convenience base?
   A good example is the claim in point 4 above where weeks 9 and 18 were chosen for analysis. There was no scientific logic in choosing these weeks. They were done to only exaggerate a claim.
c. Are we using 'commonsense'?

The answer to the first question should be "yes". That for the second question: "No". The answer to the third question should hopefully be "yes" (as in point 3 above).

11. Use the right basket for analysis

TAM does not classify channels into specific genres. But it is expected that the user includes all channels of a genre when analyzing data for that genre. This is based on her domain knowledge. Similarly the case of programmes though this is admittedly more complicated. But the full attempt should be to give an honest representation of the environment in which the channel or programme exists without prejudice to any other player.

12. <u>Beware of 'blow-up' estimates</u>

We define "blow-up estimates" as derived estimates that are computed on shrunk denominators thus exaggerating its importance in the overall data context.

Two specific examples are:

A. Relative Shares
B. Time Spent per Viewer.

Both are covered specially in rules below.

13. <u>Use and/or Interpret Relative Channel Shares with discretion</u>

A typical industry practice is to compute relative shares of channels within a genre. In such a case it is important to:

- Ensure that all channels in the genre have been taken for analysis (Point 11 above)
- State the associated base audience size clearly. A channel could command a 70% relative channel share in a genre that accounts for 0.5% of the total average audience size.
- When using relative channel shares, clearly state that these are "relative" to other channels in the genre.

Person in population	Channel A	Channel B
1	20	5
2	-	5
3	-	5
4	-	5
5	-	5
6	-	5
7	-	5
8	-	5
9	-	5
10	-	5
11	-	-
12	-	-
.....	-	-
99	-	-
100	-	-

	Channel A	Channel B
Reach	1%	10%
TSU (mins)	0.2	0.5
TSV (mins)	**20**	**5**
Share	29%	71%

14. <u>Use and/or Interpret Time Spent per Viewer (TSV) with discretion</u>

Count all the minutes that a channel is viewed in a geography. Divide it by the population estimate (viewers and non-viewers of that channel) of that geography. The result is Time Spent Per Universe (TSU). This is what TAM reports.

Now divide the total man-minutes by only the number of viewers of the channel. This is Time Spent Per Viewer (TSV). It stands to reason that TSV >= TSU.

As a hypothetical example, take the case of a market whose population is 100 individuals. Let's assume that only two channels (A and B) are available to be viewed in this small market. Their viewing behavior for a certain time period is given on the left.

Now Channel A states that "We are viewed four times more than Channel B". [1] This claim is made on the basis of TSV but hides the fact that its Reach is just 1% and therefore accounts for just 29% of the average audience size. Moreover Channel B has a very stable TSV (5 minutes across viewers).

The TSV by its nature is a very unstable estimate. It is not good practice to analyze data for a single day, for example. it was the misuse of the TSV estimate that TAM stopped reporting it in its software. However, this can still be computed by a user and some cases of misrepresentations continue.

III. Graphical representations

In many cases, graphs are a major improvement over showing data in a table. But graphs are most often more liable to misrepresentation. Default options available in commercial graphing software are not the ideal to represent data. Moreover some of them contain options that are patently bad from a scientific information display perspective. There is a lot of literature on what constitutes good data representation. Without getting too academic about this topic, some simple guidelines to be followed when representing TAM data by a graph are given below.

15. Do not use 3D graphs

16. The horizontal and vertical axes in a graph should be perpendicular to each other

Typically the data you are displaying has only one dimension or two at the most – TVR% or Channel Share for example. Typically, 3D graphs are never useful to represent data. In fact, the single purpose of using them seems to be either to mislead or to give some supposed visual appeal. Chart below combines both evils of points 15 and 16 above.

[1] As noted in point 6, the base estimate to declare viewership must be made clear. TSV, in this case.

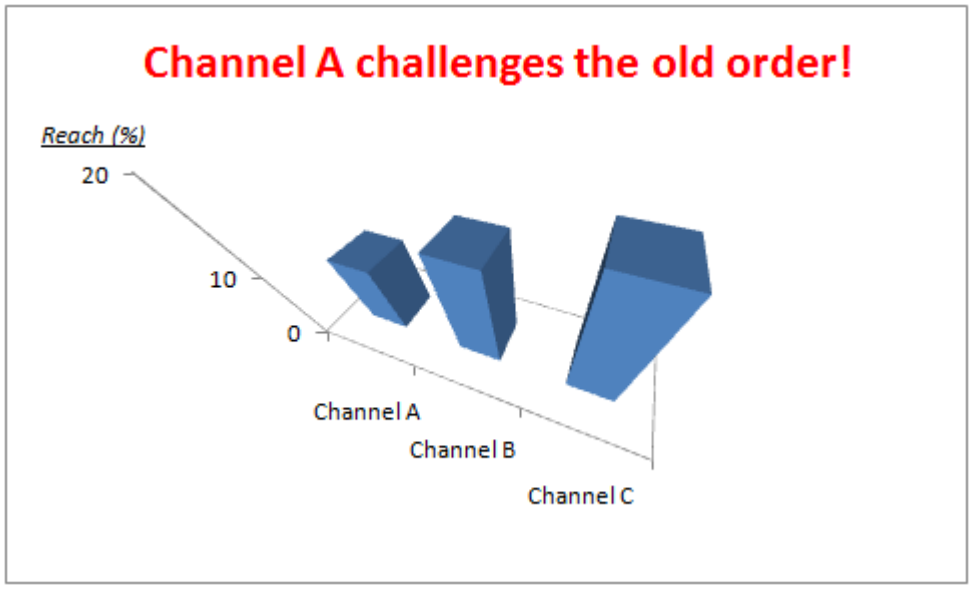

The actual table used to draw the above chart is below. Just shows how a combination of axes tilting and 3D charts have completely misrepresented the differences among the channels.

Reach (%)		
Channel A	Channel B	Channel C
10	15	18

For a starker example, see the chart below.

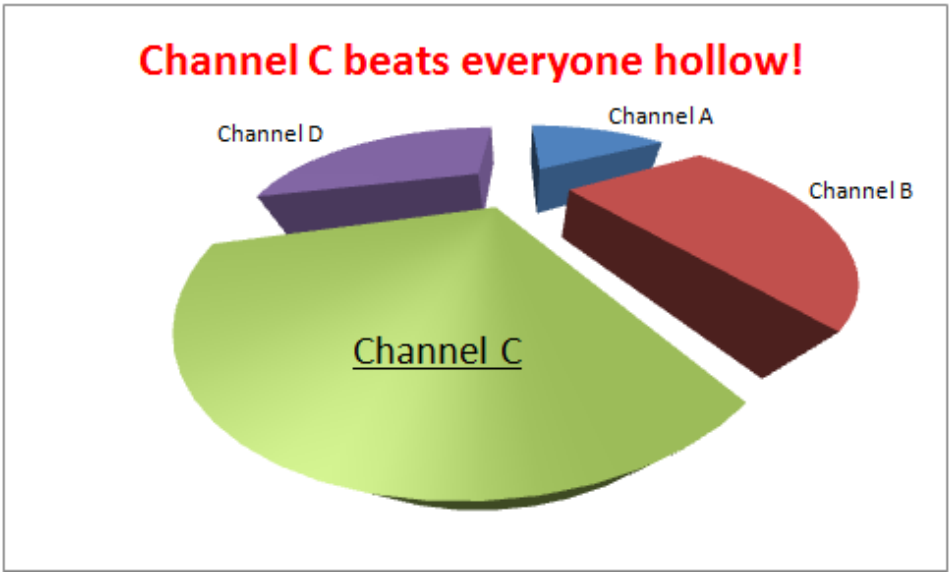

The actual data table used to draw this chart was the one below:

Channel Share (%)			
Channel A	Channel B	Channel C	Channel D
10	15	18	18

In passing we also note that pie-charts are not a very good way of displaying data since the human eye cannot compare areas (the 'pies') as well as linear distance (as in a line graph)

17. Be careful about the starting point for the vertical axis

For this, see the graph for rule 2 where the starting point of 0.17 turns out to be a misleading starting point

18. Standardize aesthetics of all data points

Rules 16 and 17 are illustrated with a chart series titled "Evolution of a chart" in the page after the next.

Evolution of a chart

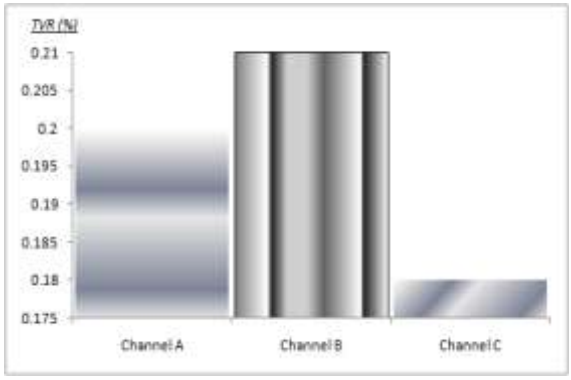

Chart 1

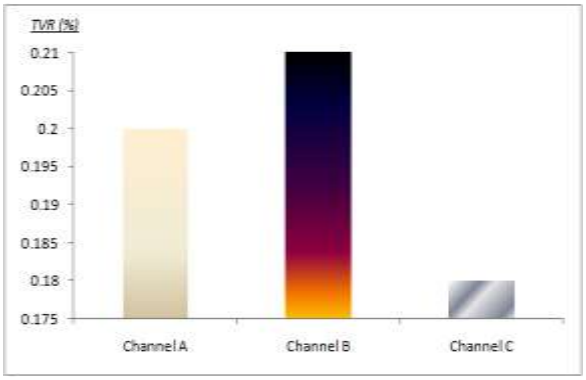

Chart 2

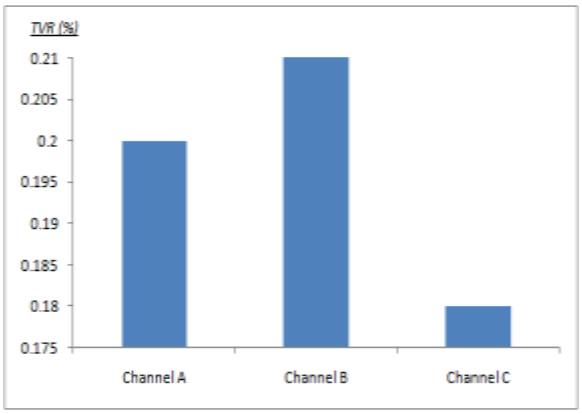

Chart 3

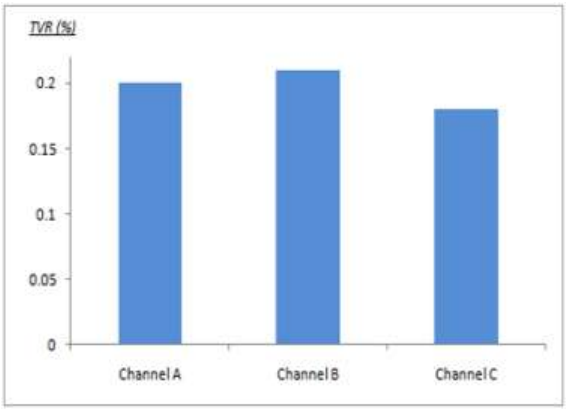

Chart 4

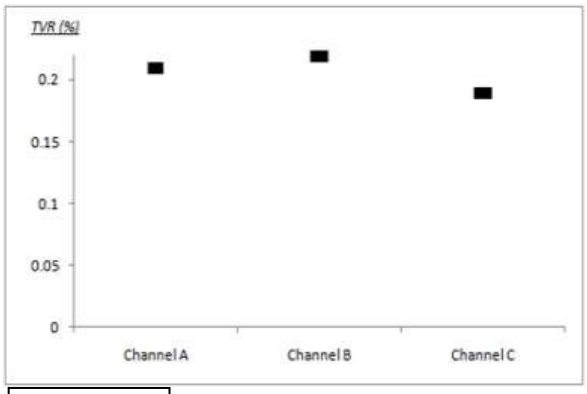

Chart 5

TVR (%)		
Channel A	Channel B	Channel C
0.2	0.21	0.18

Table 1

The first chart in the series above uses striations to distinguish the bars but is unnecessary. In fact, the patterns produce what are called Moiré vibrations which tend to create visual mirages not corresponding to the actual underlying data.

Chart 2 in the series does away with the striations but now introduces colour – once again misleading since the attempt here seems to be to use stronger colours for one's own channel.

Chart 3 standardizes the aesthetics but continues to exaggerate the differences between the channels by artificially starting the vertical axis at a non-zero base.

See how in chart 4 now all channels seem more or less equal after having based the axis properly (Rule 17).

Chart 5 is even better since it decides to waste less ink and focuses only on showing the three points on the graph.

But the best is reducing even this to just the three data points. Coming to think of it, this is all you actually need! This is not to suggest that graphs should never be used. On the contrary wherever possible they should be used to highlight the relationships in the data. But these should be done with due care to truth.

Comparing Chart 1 and Table 1 shows what a dramatic perceptual difference there is and how charts can mislead.

19. <u>A good example</u>

As an illustrative case (<u>purely illustrative</u>), the following are the parameters a reasonably good analysis will take as an input. This is for an English special content channel. Once again, note that we are not suggesting that only these parameters should be taken for analysis.

i. TG: SEC AB 25 years+
ii. Geography: All metros i.e. top 8 metros and 1 Mn.+ cities.
iii. Weeks: Last 13 weeks
iv. Dayparts: All day – weekend and weekday
v. All channels in the competitive basket taken
vi. Reach at 1 Min+ was analyses. TVR was not taken since this was found to be very small to make practical or statistical sense. On the other hand,

reach was found to yield good insight and establish some relationship with content.

vii. What was displayed were Point estimates (thirteen-week average reach) and trend. The first was displayed as a table and the second by was of a line graph with the vertical axis starting at zero.

viii. Headers and footers clearly noted the analysis parameters and were easily readable by the end consumer of the communication.

IV. End-Note

A rise or drop in viewership for a programme need not mean that TV viewers have suddenly started liking or disliking content. Analyzing TV viewing behavior is more complex than it seems. There is a range of factors at play. Environmental factors like power cuts, competitive programmes, effect of promotions, special events like cricket etc. all contribute to behavioral and therefore data changes. These must be studied carefully before jumping to conclusions. Specifically, the current digitization effort across the country had led to a fundamental change in the reception mode by a home. During his adoption phase, it is important for the user to keep a special eye out for elements that would create either enhancement or disruption in his viewing behavior. One tip for users is to maintain an event diary that changes in viewing can be associated with.

We are living in a data intensive world and the intensity will only increase. However with this intensity comes greater responsibility in using and interpreting data. We hope this was a useful set of guidelines. Based on user feedback and suggestions, we hope to enhance this note and release expanded versions in the future.

Appendix 1

As illustrated below, the Standard Error reduces as the sample sizes increase; therefore we strongly recommend –

1. **Use a demographic definition that will be based on the largest possible sample size. To achieve this, we strongly recommend users to combine markets AND desist from splitting demographics too fine**
2. **Use a minimum of 13 week averages (moving or week-to-date) AND on a target audience definition that will work off a minimum sample base of 1,000 respondents**

Standard Errors (SEs) at various estimate levels for different sample sizes

	Sample Size (Individuals)												
Reach/TVR	50	100	200	300	400	500	1000	1500	3000	5000	10000	15000	20000
0.1%	0.6%	0.4%	0.3%	0.2%	0.2%	0.2%	0.1%	0.1%	0.1%	0.1%	0.0%	0.0%	0.0%
0.2%	0.8%	0.6%	0.4%	0.3%	0.3%	0.3%	0.2%	0.2%	0.1%	0.1%	0.1%	0.0%	0.0%
0.5%	1.3%	0.9%	0.7%	0.5%	0.5%	0.4%	0.3%	0.2%	0.2%	0.1%	0.1%	0.1%	0.1%
1%	1.9%	1.3%	0.9%	0.8%	0.7%	0.6%	0.4%	0.3%	0.2%	0.2%	0.1%	0.1%	0.1%
1.5%	2.3%	1.6%	1.1%	0.9%	0.8%	0.7%	0.5%	0.4%	0.3%	0.2%	0.2%	0.1%	0.1%
2%	2.6%	1.9%	1.3%	1.1%	0.9%	0.8%	0.6%	0.5%	0.3%	0.3%	0.2%	0.2%	0.1%
2.5%	2.9%	2.1%	1.5%	1.2%	1.0%	0.9%	0.7%	0.5%	0.4%	0.3%	0.2%	0.2%	0.1%
3%	3.2%	2.3%	1.6%	1.3%	1.1%	1.0%	0.7%	0.6%	0.4%	0.3%	0.2%	0.2%	0.2%
3.5%	3.4%	2.4%	1.7%	1.4%	1.2%	1.1%	0.8%	0.6%	0.4%	0.3%	0.2%	0.2%	0.2%
4%	3.7%	2.6%	1.8%	1.5%	1.3%	1.2%	0.8%	0.7%	0.5%	0.4%	0.3%	0.2%	0.2%
4.5%	3.9%	2.7%	1.9%	1.6%	1.4%	1.2%	0.9%	0.7%	0.5%	0.4%	0.3%	0.2%	0.2%
5%	4.1%	2.9%	2.0%	1.7%	1.4%	1.3%	0.9%	0.7%	0.5%	0.4%	0.3%	0.2%	0.2%
5.5%	4.3%	3.0%	2.1%	1.7%	1.5%	1.3%	1.0%	0.8%	0.6%	0.4%	0.3%	0.2%	0.2%
6%	4.4%	3.1%	2.2%	1.8%	1.6%	1.4%	1.0%	0.8%	0.6%	0.4%	0.3%	0.3%	0.2%
6.5%	4.6%	3.3%	2.3%	1.9%	1.6%	1.5%	1.0%	0.8%	0.6%	0.5%	0.3%	0.3%	0.2%
7%	4.8%	3.4%	2.4%	1.9%	1.7%	1.5%	1.1%	0.9%	0.6%	0.5%	0.3%	0.3%	0.2%
7.5%	4.9%	3.5%	2.5%	2.0%	1.7%	1.6%	1.1%	0.9%	0.6%	0.5%	0.3%	0.3%	0.2%
8%	5.1%	3.6%	2.5%	2.1%	1.8%	1.6%	1.1%	0.9%	0.7%	0.5%	0.4%	0.3%	0.3%
8.5%	5.2%	3.7%	2.6%	2.1%	1.8%	1.6%	1.2%	1.0%	0.7%	0.5%	0.4%	0.3%	0.3%
9%	5.4%	3.8%	2.7%	2.2%	1.9%	1.7%	1.2%	1.0%	0.7%	0.5%	0.4%	0.3%	0.3%
9.5%	5.5%	3.9%	2.7%	2.2%	1.9%	1.7%	1.2%	1.0%	0.7%	0.5%	0.4%	0.3%	0.3%
10%	5.6%	4.0%	2.8%	2.3%	2.0%	1.8%	1.3%	1.0%	0.7%	0.6%	0.4%	0.3%	0.3%

1. TAM does not report data below the threshold sample size of 50 individuals

2. Based on a complex sampling scheme with a design and weighting effect of 1.75

3. Use this table as a guideline only for point analysis, not for calculating SEs of GRPs, averages etc. For example, taking averages can give upto 40% reduction in SEs

4. SEs are rounded to a single digit after the decimal point