

## ASSESSING THE IMPACT OF A BRAND CRISIS USING BIG DATA: THE CASE OF THE VW DIESEL EMISSION CRISIS

By Marco Vriens, Chad Vidden, Song Chen, University of Wisconsin, La Crosse  
and Sandro Kaulartz, Ipsos

### ABSTRACT

The study of market structure analysis and brand associations are strategic tools for marketers that can aid them in tackling brand and product positioning challenges. Often surveys are used to gain such insights. In brand or product crises cases, timing is an essential part of the insights, and surveys become inadequate to quickly assess the severity of the issues. In product launch scenarios, surveys may not be agile enough. We will show that text mining of consumer online-generated data (CGD) along with network analytics can give marketing a new set of tools to quickly assess and diagnose the impact of a brand crisis on overall market structure and on changes in brand associations and brand metrics. We use the VW Diesel Emission Crisis to illustrate these new tools.

### INTRODUCTION

The study of market structure analysis and brand associations provides strategic tools for marketers that can aid them in tackling brand and product positioning challenges. Up until now practical market structure and brand research have relied on survey research for the most part. Survey research is expensive, and once the data has been collected, the timing of the snapshot is fixed, i.e. we cannot expand the scope of the view. This is an issue because the initial design really sets the limits for what can be done analysis-wise. Any mistakes or omitted brands or SKUs will leave the researcher and their client frustrated. In addition, we cannot go back in time. Imagine the CMO of VW brand reading his morning newspaper on September 21, 2015. VW confirmed on that day to stop the sales of all four-cylinder diesel cars. What will be the real impact of this breaking news that VW knowingly and on purpose adjusted their diesel emission results? To do a survey would take too long and the CMO would still not know what would have changed from the situation prior to the crisis.

Several recently proposed approaches leverage consumer-generated data (CGD) such as discussions on product review websites, forums, blogs, etc. to study market structures and brand associations<sup>1,2,3, 4,5</sup> using a combination of text and network analytics. These approaches have two key benefits over traditional survey based approaches. First, the scope can easily be changed and can be much broader (bigger) than what would be feasible with survey research. For example, one study included more than 150 sub brands, another study included more than a 1000 products. Such scopes would not be feasible with surveys. If those studies would have started out with too few brands, they could have easily cast a wider net to scrape the online conversations. Second, online CGD has a time stamp; hence, we can collect data before, during and after a campaign, before and after a crisis, or simply over a period of recurring time intervals to see if there are changes in the market.

The contribution of our paper is twofold. First, we show how online CGD can be used to assess sudden shocks in the market. Online CGD were used by Netzer et al.<sup>1</sup> to derive insights in evaluating marketing campaigns over a 5- to 6-year period by showing how Cadillac's position moved over time. Online CGD coupled with the right analytics can similarly be used to gain insight into issues that play out over shorter periods of time. Marketing challenges such as new product introductions and brand crises, such as Toyota's recall, the Wells Fargo cross-selling scandal, and the VW Diesel Emission scandal are examples of marketing problems where a firm quickly needs to determine the impact of such shocks in the market. The timing that such events "shock" the market or shock the brand's performance is most likely less than one year. We use CGD to showcase how we can shed light on the VW Diesel Emission Crisis. Specifically,

we investigate if and how VW's position in the social market structure map significantly changed as a result of the scandal (by comparing the market structure before the crisis with the market structure after the crisis. For this purpose, we derive market structure based on co-mention data.

Second, we look at brand association maps, and we employ a new network metric to investigate if these provide additional insights into the nature and severity of the scandal. We propose and utilize a new network metric, which we call brand density. To our knowledge this type of analysis has not been done yet using online CGD.

The rest of this paper is organized as follows: The next section will briefly review the key selected papers in this area. Then we outline the methodology, data and analysis used in our study. We collected data from prior to the VW crisis, to up until August of 2016. In the results section, we show the results of a variety of approaches that have been used to analyze and model online CGD. We believe the approaches described by Netzer et al. and Gensler et al.<sup>1,3</sup> to be most useful. We employ methods not previously attempted including some new visualization techniques. We conclude with practical recommendations to make this type of analysis available to a wider audience.

### Previous Research

The use of consumer-generated data on the Internet via product reviews websites or other social media sources has been an increasingly popular topic in marketing science and marketing analytics. Researchers have focused on the extraction and processing of CGD and summarizing the data by high-level metrics such as overall positive or negative sentiment (e.g. average number of "stars" given by product reviewers<sup>6</sup>). The field sought ways to apply more statistically sophisticated analytics to such CGD. Some developments have focused on predictive insight: e.g. using CGD to predict product sales or consumer choices. For example, Archak et al.<sup>7</sup> text-mined online product reviews for cameras and camcorders and extracted the attributes reviewers use to describe the products and estimated the overall impact on sales. Others have focused on an in-depth understanding of the market structure: the degree to which brands are competing with each other, and on understanding the differences and similarities across brands in terms of the attributes and other associations, they seem to induce in online conversations. This is the focus of our paper. Only a handful of studies of this pioneering work have been done to date on the topic of using online CGC to identify market structures and brand association insights.

Lee and Bradlow<sup>2</sup> extracted data from a product review website (Epinions.com), pertaining to cameras. This particular website separated user-provided pros and cons of each product. This made the text mining simpler compared to sites where users can provide unstructured comments. The brand and products were coded in terms of what positive attributes and what negative attributes the product possessed, according to the user. Then they applied a clustering analysis of the word vectors. The reason for this was to identify higher-level product dimensions (e.g. "ease of use" could be a cluster of different concrete features that all improve the ease of use). The authors suggested three ways in which this data can yield managerial useful insights. First, the CGD and corresponding clustering of the attributes yielded new insights into what attributes and dimension may be important to consumers. They compared the attributes derived from the CGD with attributes encountered in expert buying guides (e.g. CNET, etc.), and they found the CGD yielded attributes not found in such guides. Second, they used a consumer survey to assess the importance of the CGD derived attributes, and they confirmed that such attributes are indeed rated as important to consumers. We note that though this is a good finding, it is not feasible to generalize such a finding. One of the benefits of using online CGD is we can avoid doing surveys for market structure and brand association insights. Third, they used correspondence analysis to obtain a two-dimensional market structure map. Using the attribute dimensions and linking these back to the brand, they created a brand attribute dimension matrix. For each attribute dimension, they counted the number of times a given brand is associated with that attribute dimension. This two-mode, two-way matrix was analyzed using correspondence analysis. In Lee and Bradlow's study, no attempt was made to give more in-depth brand insight by analyzing the brand association in more detail; furthermore, no attempt was made to utilize the timing component of online data -- the ability to see how things change over time.



## 2017 DMA Analytics Journal

Netzer et al.<sup>1</sup> used CGD to text-mine discussions about cars extracted from the Edmunds.com website. Their text-mining approach consisted of a combination of supervised machine learning and rule-based text mining. The text-mining was used first to derive the so-called brand co-mention matrix. The frequency of co-mention across a set of product reviews or online discussions was used as a raw indicator of the similarity between two brands. Their data was extracted on 169 car sub-brands (e.g. Honda Civic, Mazda 6, etc.), and 30 car parent brands (e.g. Honda, BMW, etc.). The raw number of co-mentions was adjusted for the total number of mentions. Popular brands would have a high number of mentions, resulting in a higher frequency of co-mentions. In order to use the co-mention metric as a valid measure of similarity, the co-mention metric was adjusted using the lift method:

**Lift (A, B) =  $P(A,B)/\{P(A)*P(B)\}$** , where P(A, B) the probability is that brand A and brand B appear together in a message or “text chunk”, whichever is being used as the unit of analysis. P(A) is the probability that brand A is being mentioned in a message, likewise P(B) is the probability that brand B is being mentioned.

After these adjustments, the matrix of adjusted co-mention measures can be analyzed with standard multi-dimensional scaling (MDS) or with spring-embedded networks (SEN)<sup>8</sup>, specifically the Kamada and Kawai approach<sup>9</sup>. MDS does have some drawbacks such as a circular positioning (the horseshoe) of the alternatives which is a sign of a degenerative solution. When the number of competing alternatives grows large, the risk seems especially likely to increase<sup>10,11</sup>. There are other visualization methods but these are beyond the scope of our paper. Netzer et al.<sup>1</sup> compared their social market structure map with market structure maps derived from survey research and one derived from actual brand switching data. The correlation between their market structure map and the brand switching market structure map was high (.79), though the correlation with the survey-based market structure map was modest (0.55). This showed that a market structure map derived from Big Data could be a valid representation of the actual market structure. They also looked more in-depth into what attributes are associated with the brand, and also found associations that were not anticipated by the brand. For example, Nissan was associated with “college”, and this seemed to indicate that this car model was being viewed as a good choice for college bound children. The authors did also look at how the brand’s positioning changed over time. They showed that over an approximate 5-year period, the lift of Cadillac over luxury brands (indicating the degree to which Cadillac is being co-mentioned with luxury brands) went up and surpassed the lift of Cadillac over American brands. This inflection happened before there was an inflection in sales.

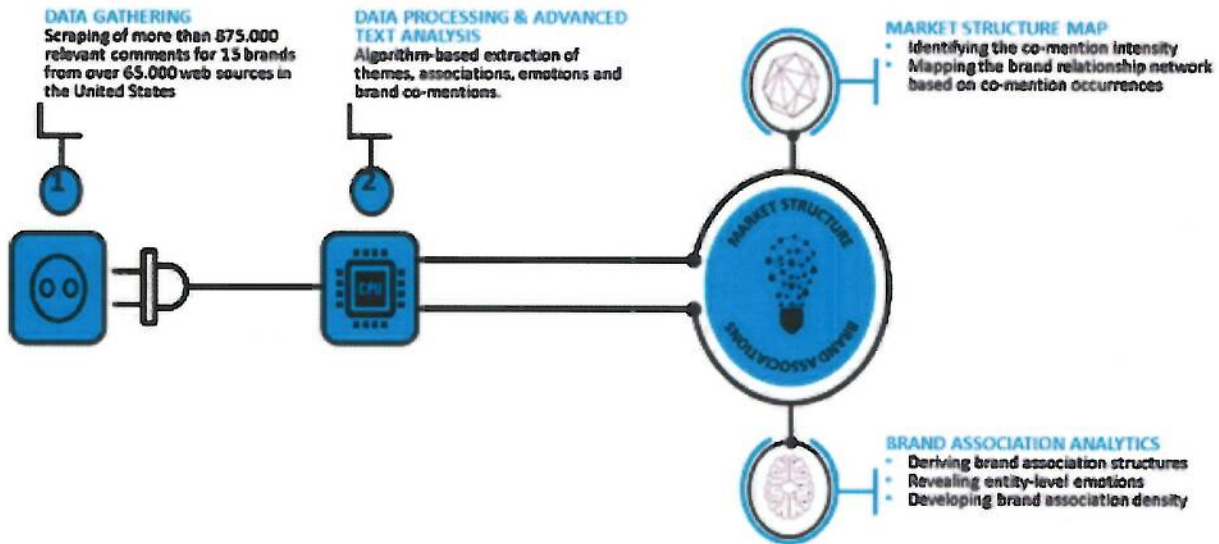
Ringel and Skiera<sup>4</sup> analyzed individual search log data to determine brand consideration sets using a price comparison website. They identified size and composition of the consideration set. From that, they derived the frequency that any two brands appeared in the same consideration set. An adjustment needed to be made to these raw frequencies. In this case, they used Linden’s normalization procedure ( $\text{Frequency}(A, B) / \{\sqrt{A}*\sqrt{B}\}$ ). The resulting matrix of adjusted relative frequencies was, in essence, a similarity matrix and could be analyzed with a variety of multivariate techniques. They used Hierarchical (Ward’s) clustering and MDS (PROXSCAL) on their data to get the market structure maps. In this study, no validation was provided. Ringel and Skiera<sup>5</sup> tackled the challenge how to gain competitive market structure insights when the number of products is very high (i.e. > 1,000). Although this is an impressive accomplishment, we do believe that most markets involve a much lower number of competitive sets.

Online CGD can also be used to extract brand insights, specifically associations with the brand that have been found to be important predictors of brand success<sup>12</sup>. The identification of brand association maps used mostly product review websites to date. Gensler et al.<sup>3</sup> used SEN not to visualize the position of brands in a map but to visualize the relationship between brands and products and brand associations.

**Data Collection Methodology**

In the following section, we illustrate our data gathering, processing and advanced text analytic process flow. The process of our online CGD data collection is shown in Figure 1.

**Figure 1:**  
Data gathering process



Web scraping technologies usually operate on Boolean language systematics to scrape CGD in the form of posts or comments related to a specific research object. The development of Boolean queries within the search refinement stage consists of terms that best characterize the relevant CGD and rule out unwanted content. Based on the Boolean query definition web crawlers then extract the CGD from URLs based on a predefined set of web sources that allows for a decent coverage of the web in a particular market. This approach often relies on a limited source set that is indexed within the crawling scope of the technology.

For our study, we used a web crawling approach that mirrors the user experience from search engines. This approach allowed us to capture all websites with relevant, publicly available CGD to search engines. This led to a maximum coverage while simultaneously returning only relevant results, as only content that can be retrieved by search engines can also be found by an information-seeking user - and potentially influence them. The crawling process within webpages (e.g. articles within threads) also ensured the removal of irrelevant content for the research purpose such as navigation content, ads, and polls on the webpage. A duplicate detection system removes copied and duplicate posts from the dataset ensuring the same content is only counted once.

The web contains CGD in private contexts, but also professional and editorial content. Both kinds are often mixed on a single webpage, such as user comments next to a product description. In order to really measure the voice of the consumer, it is very important to be able to properly differentiate between user-generated and editorial content. We applied an advanced retrieval method to automatically distinguish between user-generated and professional content, delivering only the content that is actually relevant for the purpose of this study. The overall collected and analysed user-generated content universe for the illustrated study comprised more than 875,000 post and comments associated with the top 15 automotive brands in terms of market share from over 65,000 different web sources in the United States.

## 2017 DMA Analytics Journal

After the data gathering stage, we deployed sophisticated language processing algorithms to identify the key themes and co-mentioned brand within each user-generated content unit. The used algorithms allowed for a bottom-up analytic approach that enabled an automatic discovery of the key themes and consumer associations within a comment or post. We applied a rule-based automated text analysis approach to deconstruct each sentence from the relevant user-generated content universe into its components to identify themes (i.e. engine), brand associations (i.e. price), and specific brand references (i.e. Volkswagen). Furthermore, we discovered the drivers and emotions (adjectives and verbs) that were expressed in direct relationship to the themes, brand associations and brands (i.e. good engine, high price, amazing Volkswagen). Being able to analyse the sentiment and emotional dimension (e.g., positive, negative, or neutral) on the entity level is indispensable to not only identify which themes and brand associations are perceived as being important but also to develop a differentiated valence analysis of various entities within one comment.

The entity-based sentiment analysis is a critical component of the dynamic market structure approach to capture the perception polarisations and interrelation between heterogeneous brand association structures (i.e. fair price but low comfort) within one comment. The described approach allowed us to explore and isolate brand co-mentions within comments as well as the mental brand association structure for all brands.

### Analyses

First, we looked at the number of mentions for VW versus the average of the other brands. Second, we compared the before/after lift metrics between VW and those other brands. The lift numbers can be viewed as a bivariate statistic as to how similarity between brands has increased/decreased. Third, we calculated before/after two-dimensional maps using MDS, SEN and a somewhat new approach called Visualization of Similarities (VOS<sup>13</sup>). The VOS approach is a special case of MDS with a specific choice of weights and dissimilarities that are derived from co-occurrences (The Smacof package in R or Proxscal in SPSS can be used for this) allowing for specifying weights per residual. These techniques should give fairly similar outcomes for the studies with a scope of, say, less than a 100 products/brands. We could also look at three-dimensional maps; but these are much harder to interpret and use by managers, so we limited ourselves to 2-dimensional solutions. Hierarchical clustering in practical situations sometimes gives results that are easier to interpret. However, this analysis did not reveal any new insights, so we left it out. Fourth, we looked at sentiment changes. This aspect has not been captured in the previous studies on this topic. Fifth, we calculated and created two-dimension brand-associations maps. For the brand-association maps, we calculated a so-called "brand density metric" to assess whether or not overall brand equity had gone up or down for VW after the crisis (following Vriens et al., 2016)<sup>14</sup>.

### Results

The results of the raw mentions and the lift numbers (before/after) are shown in Figures 2 and 3.



# 2017 DMA Analytics Journal

Figure 2:  
Mention Percentages by Month





Figure 3:  
Change in VW Lift, After-Before Crisis




## 2017 DMA Analytics Journal

Figure 2 shows that there was a sharp increase in the raw number of mentions of VW which lasted about a quarter and remained somewhat elevated to where it was before the crisis. Based on this data, it seems that word-of-mouth for VW was increasing; but it is unclear if the content of these discussions had substantially changed. Figure 3 gives a first clue as to what was happening and shows that VW had become less similar to most brands with the exception of Chrysler, Buick, and Chevrolet. These American brands clearly still have brand image problems to fix, so this movement was not a good one for VW.


Figures 4a and 4b show the before/after results of a metric MDS solution (4a) and a VOS solution (4B) on the lift matrix. We are not showing the hierarchical clustering results as the insights were fairly comparable to the MDS solutions.

**Figure 4a:**  
**Metric MDS Before and After Crisis**



## 2017 DMA Analytics Journal

Figure 4b:  
VOS Before and After Crisis



Both metric and VOS in our case were fairly similar, though the VOS solution was somewhat easier to interpret as we saw VW (and Audi to a lesser degree) move away from BMW and away from the Asian brands toward the American brands. We also ran SEN analyses where the results were very similar to the MDS so we left these out here.

Both MDS solutions show only a fairly marginal movement of the VW brand. Before the crisis, VW seemed to have an admirable position right between the luxury brands of Audi and BMW and the reliable Asian brands. After the crisis, VW seemed to fall between Audi and the American brands, which we believe is less desirable. We note that Audi seemed to move with VW as well, and it may be hard for VW to be completely de-coupled from Audi as many consumers will know they have the same parent company. Consequently, Audi would likely be affected by VW's crisis.




## 2017 DMA Analytics Journal

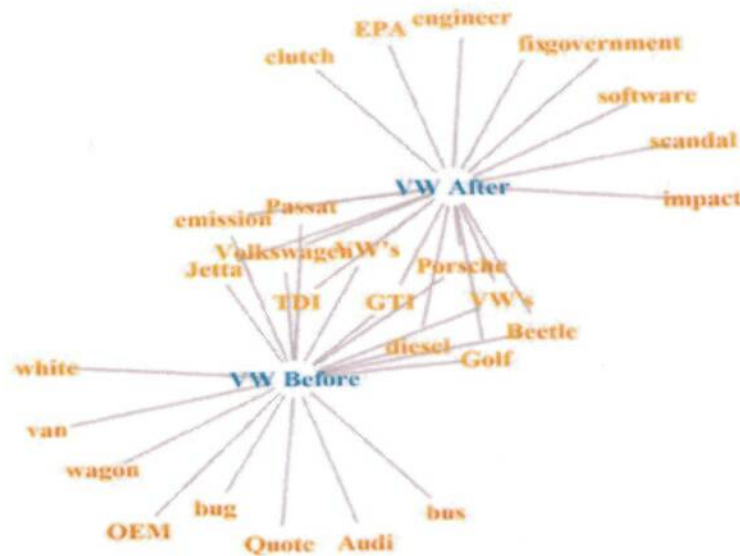
Next we looked at sentiment changes. Figure 5 shows how the sentiment around VW changed. The negative and emotional comments increased dramatically after the crisis announcement. We note that the category “emotional” comments is for the most part the sum of positive and negative comments but there is a small part of the comments that were labeled “emotional” without a further positive or negative qualification next to it.

The brand association maps were also revealing (see Figure 6), whereas before the crisis VW had mostly positive and neutral associations. Afterward, it received much more negative ones, and we saw that what came first to mind were the topics related to the scandal.

**Figure 5:**  
**Average Change in Brand Sentiment**



**Figure 6:**  
**Volkswagen Crisis Concept Map**



## 2017 DMA Analytics Journal

We also calculated a density metric. The nature of online CGD is different than individual level survey data, so we had to amend the calculation of the metric a little bit. Also, calculating an overall brand density metric in this case did not give us any insight; because it was not the raw number of associations per se that was changing, but the types of associations that were changing. The results are shown in Figure 7.

Figure 7:  
Positive and Negative Brand Density (Before/After)

	POSITIVE DENSITY RANK		NEGATIVE DENSITY RANK	
	BEFORE	AFTER	BEFORE	AFTER
Audi	1	1	4	2
BMW	7	4	6	7
Buick	15	15	15	15
Chevrolet	9	11	11	11
Chrysler	12	12	12	12
Dodge	4	6	1	3
Ford	11	10	9	9
Honda	5	2	5	6
Hyundai	14	14	13	13
Kia	13	13	14	14
Mazda	10	8	10	10
Nissan	6	5	8	4
Subaru	8	9	7	8
Toyota	3	3	2	5
Volkswag	2	7	3	1

The results were dramatic. From a positive brand density perspective, VW dropped from a rank of 2 to a rank of 7. The negative brand density jumped from 3 to 1 meaning VW now had the highest number of negative associations. Applying the Zipf transformation<sup>15</sup> - to gauge what drop in market share VW could expect - showed a jaw-dropping 12 share point (dropping from 17.5 to 5.5).

### Conclusion

In this paper, we showed how online consumer generated data (CGD) can be used to quickly and cost-efficiently assess the potential impact of a brand or product crisis. Such crises, especially with the quick dissemination of news through the Internet, are events firms really have to worry about as they can dramatically affect company performance both in the short and the long run. Think about the Wells Fargo scandal (2016), the Takata airbag crisis (2016), the GM faulty ignition switch recall crisis (2014), the Blue Bell tainted ice cream crisis (2015), etc.

We used CGD in combination with network analytics and network visualization approaches to assess and diagnose how VW suffered from its Diesel Emission crisis. We showed that VW moved away from BMW and the Asian brands while moving toward American brands. We used several network visualization techniques (MDS, VOS, and SEN), all of which more or less produced the same insights. In the context of online CGD, some authors have argued that SEN, and other new techniques such as VOS may be better suited with increasing numbers of products or brands to be mapped. There is some evidence that MDS's stress value deteriorates with an increasing number of stimuli to be mapped<sup>10</sup>.

## 2017 DMA Analytics Journal

We also mapped the brand associations; and we showed how associations for VW before the crisis were very different than after the crisis; and the number of negative and emotional comments that were used to talk about VW goes up dramatically after the crisis. Lastly, we calculated a new network metric: the brand density metric and showed that positive brand density goes down whereas negative brand density goes up. This likely will have significant effects on VW's market share.

### Acknowledgements

We would like to thank Oded Netzer (Columbia University, New York) for his advice on the data collection and the analysis. We also would like to thank Frederick Borchers and Stefan Maas from Ipsos for their help in managing the data collection process. Finally, we would like to thank Jeff Ross (University of Wisconsin-La Crosse) for proofreading our paper multiple times.

### References

1. Netzer, O., R. Feldman, J. Goldenberg, and M. Fresko (2012), "Mine Your Own Business: Market Structure Surveillance Through Text Mining", *Marketing Science*, 31, 3, 521-543.
2. Lee, T. Y., E.T. Bradlow (2011), "Automated Marketing Research Using Online Customer Reviews", *Journal of Marketing Research*, Vol. XLVIII, 881-894.
3. Gensler, S. F. Voickner, M. Egger, K. Fishbach, D. Schoder (2016), "Listen to Your Customers: Insights into Brand Image Using Online Consumer-Generated Product Reviews", *International Journal of Electronic Commerce*, 20, 1, pp. 112-141.
4. Ringel, D. M. and B. Skiera (2014), "Understanding Competition Using Big Consumer Search Data", 2014 47th Hawaii International Conference on System Science, pp. 3129-3138.
5. Ringel, D. M. and B. Skiera (2016), "Visualizing Asymmetric Competition Among More than 1000 Products Using Big Data Search", *Marketing Science*, 35, 1, 511-534.
6. Chevalier, J. A. and D. Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews", *Journal of Marketing Research*, 43 (3) pp. 345-354.
7. Archak, N. A. Ghose, P. Ipeirotis (2011), "Deriving Pricing Power of Product Features by Mining Consumer Reviews", *Management Science*, 57, 8, pp. 1485-1509.
8. Eades, P. (1984), "A Heuristic for Graph Drawing", *Congressus Numerantium*, 42, 11, 149-160.
9. Kamada, T. and S. Kawai (1989), "An Algorithm for Drawing Generalized Undirected Graphs", *Information Processing Letters*, 31, 1, pp. 7-15.
10. Borg, I., P.J. F. Groenen (2005), *Modern Multidimensional Scaling: Theory and Applications*, Springer.
11. Buja, A., D.F. Swayne, M.L. Littman, N. Dean, H. Hoffman, and L. Chen (2008), "Data Visualization with Multidimensional Scaling", *Journal of Computational and Graphical Statistics*, 17, 2, 444-472.
12. Krishnan (2006), Characteristics of memory associations: A consumer-based brand equity perspective, *International Journal of Research in Marketing*, 13, 389-405.
13. Van Eck, N.J., and L. Waltman (2007), "VOS: A New Method for Visualizing Similarities Between Objects", Springer Berlin.
14. Vriens, M. C. Ho, S. Chen, C. Vidden (2016), "Predicting market share using brand density metrics: measurement, reliability and validity", Working paper.
15. Hofmeyr, J., V. Goodall, M. Bongers, and P. Holzman (2008), A New Measure of Brand Attitudinal Equity Based on the Zipf Distribution, *International Journal of Market Research*, 50, 2, 181-202.