A MEMETIC ALGORITHM FOR MAXIMIZING EARNED ATTENTION IN SOCIAL MEDIA

Pedro Godinho*
CeBER and Faculty of Economics, University of Coimbra
Av. Dias da Silva, 165, 3004-512 Coimbra (Portugal)
Phone: +351 239 790571
E-mail: pgodinho@fe.uc.pt

Luiz Moutinho
Professor of BioMarketing and Futures Research
DCU Business School
Dublin City University, Collins Avenue, Dublin 9 (Ireland)
E-mail: luiz.moutinho@dcu.ie

Margherita Pagani
Department Markets and Innovation
EMLyon Business School
23 Avenue Guy de Collongue - 69134 Ecully Cedex (France)
Phone: +33 04 78337936
E-mail: pagani@em-lyon.com

This is a post-peer-review, pre-copyedit version of an article published in Journal of Modelling in Management:
Pedro Godinho, Luiz Moutinho, Margherita Pagani, (2017) "A memetic algorithm for maximizing earned attention in social media", Journal of Modelling in Management, Vol. 12 Issue: 3, pp.364-385, https://doi.org/10.1108/JM2-10-2015-0078
The final version of the article is available online at: https://doi.org/10.1108/JM2-10-2015-0078

* Corresponding author

A MEMETIC ALGORITHM FOR OPTIMIZING EARNED ATTENTION IN SOCIAL MEDIA

Abstract

With the advent of social media in our lives and the transformation of consumer behaviour through the impact of Internet Technology, online brand-human interactions are crucial in the consumer decision-making process, as well as on corporate performance. This study develops a model to predict behavioural brand engagement as measured in terms of the amount of consumer's earned attention. The exogenous variables adopted in the model comprise longitudinal behavioural parameters related to online traffic, flow of consumer-initiated brand commentaries and the quantity of brand mentions. To test and validate the research model, we apply a Memetic Algorithm (MA) which is well tailored to the phenomenon of propagation and social contagion. This evolutionary algorithm is assessed through the comparison with a standard alternative procedure – the Steepest Ascent (SA) heuristic. Results show that the shape of the utility functions considered in the model has a huge impact on the characteristics of the best strategies, with actions focused on increasing a single variable being preferred in case of constant marginal utility, and more balanced strategies having a better performance in the case of decreasing marginal utility. Insights and implications for research and practice are then provided.

Keywords: Social networks; Memetic algorithms; Optimization; Word-of-mouth; Brand engagement; Earned attention

1. INTRODUCTION

As social technologies allow customers to connect to each other to provide information to influence the customer's buying decision, traditional retail and online companies are more and more interested in considering online videos and social networks as emergent tools which help them to generate website traffic, build brand buzz, collect inbound links and extend social media reach. The main problem that companies are facing is to gain user's attention in social media in order to increase user engagement and which social media tools to implement in order to achieve goals in the short and long term planning horizon.

In the new social media-driven business model defined by customer connectivity and interactivity, content goes hand in hand with technology, producing far-reaching effects for the way marketers influence current and potential customers (Hanna et al. 2011). All of these social technologies allow consumers to freely discuss the brand on the brand's own website, conduct product and price comparisons, and thus/consequently encourage consumer engagement with the brand (Dennison et al. 2009; Kastelein and Rempt 2010).

Moreover, consumer attention is fragmented across a plethora of contents distributed through smartphones, tablets, and desktops and consumer behavior is impacted by how and where contents are accessed. That's why it's more and more important for retailers and online brands to maximize earned attention and consequently increase brand engagement leveraging on an integrated digital marketing strategy which includes social media and new interactive digital technologies (video on the web, mobile devices and TV).

Social Network Sites (SNSs) represent an ideal tool for electronic Word-Of-Mouth (eWOM), as consumers freely create and disseminate brand-related information in their established social networks composed by friends, classmates and other acquaintances (Vollmer and Precourt 2008). Through these interactions, consumers voluntarily display their brand preference along with their persona (e.g. name and picture), which can engender eWOM communication.

Prior studies (Godwin, 1994; Kim, 2000; Koh et al., 2007; Williams and Cothrel, 2000) propose a set of stimulation drivers encouraging participation in virtual social networks such as the level of leadership involvement, the presence of offline interaction, usefulness, and IT infrastructure quality. Current research in this area is looking at different factors explaining the participation in virtual social networks. There are network effects (online social networks become more useful as more users join), critical mass theory (a sufficient number of adopters), form of the production function (how social media services focus on some or all of the different functional building blocks described), role of social referral (traffic to a site via content and/or links posted on social outposts). Previous studies also investigated behavioral contagion (Marsden 1998) defined as the spread of behaviors through populations by simple exposure. We consider in this study the phenomenon of eWOM as a specific case of social contagion and explore, using the nascent discipline of memetics, whether this established field of social science can influence online earned attention and drive brand engagement.

Memetic Algorithms (MAs) are a hot topic nowadays, mainly due to their success in solving many hard optimization problems. Unlike traditional Evolutionary Computation (EC) methods, MAs are intrinsically concerned with exploiting all available knowledge about the problem under study; this is something that was neglected in Evolutionary Algorithms (EAs) for a long time.

The exploitation of problem-specific knowledge can be accomplished in MAs by incorporating heuristics, approximation algorithms, local search techniques, specialized recombination operators, truncated exact methods, etc. Also, an important factor is the use of adequate representations of the problem being tackled.

The research model revolves around a fitness function that attempts to maximise/optimise the emerging behavioural construct of "earned attention", taking into account the explanatory role of a number of key parameters. Two of the exogenous variables are purely related to online consumer behaviour – page impressions and the total number of visitors to the website. Three other independent measures refer to the "conversation economy" and the flow of commentaries in social media, as quantitatively defined by specific time frames. Finally, the last two manipulated variables pertain to the realm of brand relationships – quantity of commentary about the brand and the number of brand mentions.

Our central argument is that social contagion (through social media and eWOM) can strongly influence and play a key role in the design, "contours" and development of brand engagement in this contemporary age and rapidly evolving environment. Our research contribution is centred around the use and application of MAs, as a computational procedure within the grouping of EAs and very similar to genetic algorithms. Still, MAs have the strength of being implicitly focused on the exploration of all available knowledge pods related to social propagation, imitation, replication and transmission of memes.

Our work draws on Godinho et al. (2015), but we extend that work in two important directions. First, we generalise the model by consider a utility-based definition of earned attention. This definition allows us to incorporate important features absent from Godinho et al. (2015), like decreasing marginal utility of the improvement in variables' values, which may lead to more balanced solutions, and different characteristics in the attractiveness of different variables. In fact, in Section 5 we will show that the use of different parameters in the utility functions has an important impact in the obtained solutions. Of course, this change has important some important impacts on the results. Second, we use an improved solution procedure. Among the innovations, we stress the use of different mutation rates, allowing the procedure to focus on convergence when the population is diverse, and in increasing population diversity when there is low diversity.

2. THEORETICAL BACKGROUND

2.1 Brand engagement

Although rooted in the disciplines of psychology and organizational behavior the concept of engagement is of increasing interest in the academic marketing literature. Prahalad & Ramaswamy (2004) were first to emphasize the importance of active customer participation and the construction of unique brand experiences.

It has been argued that engagement has a direct impact on the intention to purchase, commitment, customer loyalty and return on investment (Anderson & Mittal, 2000; Heskett et al., 1994; Porter et al., 2011; Reichheld, 2003; Reichheld & Sasser 1990) Customer engagement better describes the experience of a customer brand relationship and serves as a better predictor of consumer behavior.

In marketing it is commonly referred to as 'customer engagement' (Bowden, 2009) defined as "a psychological process that models the underlying mechanisms by which customer loyalty forms for new customers of a service brand as well as the mechanisms by which loyalty may be maintained for repeat purchase customers of a service brand" (Bowden, 2009 p. 65). In his model Bowden (2009) proposes customer engagement as a process which includes: (1) the formation of a state of calculative commitment for new customers which is considered to be a largely cognitive basis for purchase; (2) Increased levels of involvement concomitantly supported by increased levels of trust for repeat purchase customers, and (3) The development of affective commitment toward the service brand which is considered to be a more emotive basis for purchase and which may ultimately eventuate in a state of enduring brand loyalty.

Literature has typically discussed engagement in the context of customer experience (Johnson and Mathews 1997; Patterson 2000; Pagani and Mirabello 2011), customer familiarity (Soderlund 2002), customer expertise (Alba and Hutchinson 1987; Matilla and Wirtz 2002), and cognitive knowledge structures (Matilla and Wirtz 2002; Moreau, Lehmann, and Markman 2001). Other studies (Hollebeek, 2011a) refer to customer engagement as a customer's individual engagement with a brand, product or organization. Within the field of marketing six major forms of engagement have been established: customer, consumer, user, brand, advertising and media (Bowden, 2009; Gambetti & Graffigna, 2010; Heath, 2009; Mersey, Malthouse, & Calder, 2010; Neff, 2007; O'Brien & Toms, 2010; Tripathy 2014). According to Gambetti & Graffigna (2010) customer, consumer and user engagement are concerned with individuals being engaged by a brand, advertisement or communication medium (e.g.: a website).

Brand, advertising and media engagement, on the other hand, adopt the opposite approach. These three types of engagement revolve around brands, advertising and media and how they stimulate engagement within the consumer. The terms brand, customer and user engagement became of specific interest to researchers with the rise of social media marketing (Brodie et al. 2011). Since then customer engagement has been a buzzword closely related to the Web 2.0, i.e. blogs, wikis, social networks and media sharing sites. Nevertheless, there is still a broad range of interpretations of engagement (Sashi, 2012) and neither practitioners nor scholars agree on one definition (Mersey et al., 2010).

From a cognitive perspective customer engagement is considered a positive state of mind that is characterized by high energy, commitment and loyalty towards a firm (Porter et al., 2011). From a different viewpoint engagement can be seen as a category of behaviors that reflects the customer's willingness to participate and cooperate with others and the firm (Porter et al., 2011). Several researchers have made reference to the existence of focal, two-way interactions between subject and object (Handelsman et al. 2005). The customer assumes the role of the engagement subject (Barnatt, 2001; Bowden, 2009), whereas brands, products and organizations represent focal engagement objects (Barnatt, 2001; Hollebeek, 2011a). Gambetti & Graffigna (2010) have identified three types of factors which emphasize the central role of costumer engagement in building and maintaining customer-brand relationships: customer-related, media-related and company-related factors.

Customer-related factors include the consumer's cognitive and emotional need to elicit pleasure and fun by purchasing products of symbolic value (Holbrook & Hirschman, 1982; Gambetti & Graffigna, 2010). In addition, the consumer perceives an unprecedented desire to play a more active role in the consumption process and co-create brands, products and services. In a broader sense this illustrates self-expression through the creation of unique and memorable social experiences (Csikszentmihalyi, 1990; Gambetti & Graffigna, 2010).

Media-related factors are related to the rapid changes and the spread of digital technologies resulting in cross-sector convergence (Gambetti & Graffigna, 2010). A particular focus is set on new ways of engaging the customer through interaction, participation and innovative creativity.

Company-related factors correspond to changes in corporate strategies in engaging the customer and foster conversational relationships with consumers using digital media (social media, web 2.0, mobile).

2.2 The behavioral perspective

Several research streams for investigating the customer engagement were identified in the academic literature (Brodie et al. 2011): a) *behavioral perspective* – concerned with the behavioral aspects of customer engagement and customers' activities, most often beyond purchase; b) *psychological* (*cognitive and affective) perspective* -investigating customers' cognitive and affective processes which are antecedents to engagement or to interaction with brands; c) *multidimensional perspective* - unifying different dimensions of customer engagement and proposing multidimensional approach; d) *social perspective* - investigating the social and network component of the phenomenon.

In this study we adopt the behavioral perspective and we develop a model to predict behavioural brand engagement as measured in terms of the amount of consumer's earned attention. The exogenous variables adopted in the model comprise longitudinal behavioural parameters related to online traffic, flow of consumer-initiated brand commentaries and the quantity of brand mentions. We position our work in the stream of studies (Sashi, 2012; Gummerus et al., 2012; Ahuja & Medury, 2010; Ojiako et al., 2012; Javornik and Mandelli 2012) which consider customer engagement as "…*behaviors (that) go beyond transactions and may be specifically defined as customer behavioral manifestations that have a brand or a firm focus, beyond purchase, resulting from motivational drivers*" (Van Doorn et al. 2011). These studies investigate consumers' activities such as word-of-mouth, recommendations, repeat purchases, social media activities, community participation, interactions with brands and similar. Such approach is often favored by practitioners as it can lead to quantification and measurement of consumers' activities in their interactions with brands.

2.3 eWOM and Social Contagion

When marketers want to reach users of SNSs, they have two choices: buy advertising or start a viral campaign (Gilbert 2009). Sunil Gupta (interviewed by Gilbert 2009) suggests that a viral campaign may be very effective, but first it's important to understand both who influences purchase decisions in online communities and which groups of users can be influenced. In the words of Gupta: *"Viral campaigns truly leverage the network aspect of these* *social networking sites. By understanding the social network of users, firms can better understand and influence consumers' behavior"*.

Social networks make available a wealth of Word-Of-Mouth (WOM) information regarding the brands and products (Bolotaeva and Cata 2011). WOM consists of exchanging marketing information among consumers and plays a fundamental role in changing consumer attitudes and behaviour towards products and services (Katz and Lazarsfeld 1955).

The importance of WOM in influencing decision making has been well accepted in marketing and advertising literature (Engel et. al. 1969; Gilly et al. 1998). Consumers tend to rely on WOM when they search information, since it is delivered by a more trust-worthy source than company-generated messages (Feick and Price 1987). Direct observation of traditional WOM is difficult, since the information is exchanged in private conversations (Godes and Mayzlin 2004) but the emergence of Internet-based media has provided a trace of electronic Word-Of-Mouth (eWOM, also called "word of mouse") exchanged in Internet communications (Hung and Li 2007). Hennig-Thurau et al. (2004) defined eWOM as "*any positive or negative statement made by potential, actual or former customers about a product or company, which is made available to a multitude of people and institutions via Internet*". Godes and Mayzlin (2004) found out strong indications that people make offline decisions based on online information and that online communications may be a proxy for offline conversations, meaning that eWOM is relevant both for online and offline buying decisions. Sonnier et al. (2011) report a significant effect of online communications on daily sales performance. The authors use automated sentiment analysis to classify online comments into positive, negative and neutral and they conclude that positive and neutral comments tend to result in increased revenues, and negative comments tend to result in decreased revenues.

SNSs represent an ideal tool for eWOM, as consumers are allowed to freely create and propagate brand-related information in their established social networks (Vollmer and Precourt

2008). Advertising on Social Networking sites allows consumers to engage in some kind of social interactions related to the brands, leading them to voluntarily reveal their brand preference along with their personality. The growing importance that eWOM behaviours in SNSs are acquiring to brand image and product sales is making the online marketing strategies more and more relevant. In particular, building brand awareness on SNSs is an increasingly important element of the business model (Weston 2008; Bolotaeva and Cata 2011).

3. MODEL FORMULATION

The aim of this study is to propose a measure for earned attention and a model and procedure for the maximization of earned attention by a company over a period of time. We use a set of variables based on Godinho et al. (2015), but instead of a simple weighted sum we use a utility-based approach for defining earned attention.

We assume that different types of marketing actions may be undertaken in order to maximize the utility associated with earned attention (for simplicity, we just refer to earned attention), for example advertising campaigns, upgrading the website, changing product design, etc. Each of these actions has a cost, and there is a budget that cannot be exceeded.

Three categories of variables are considered in the model: 1) variables related to the propagation of brand-related information among constituents of a social network (Vollmer and Precourt 2008), with the flow of commentaries being framed in three different time periods; 2) brand mention counts and quantity of commentary; and 3) variables related to online traffic - impressions and visitors.

We assume that earned attention is positively related to the number of page impressions per different users, quantity of commentary, brand mentions and number of commentaries sent to friends (framed in three different time periods), all these as a fraction of the number of website visitors, and also the number of visitors. We assume that the marginal utility of increases in the values of these variables does not have to be constant – we expect it to be decreasing in most cases. Finally we consider that the importance given to different variables may change from company to company, so we consider different weights that can be calibrated according to the company under consideration.

Using Y_t to denote the earned attention measure at period t, we define

$$Y_t = \gamma_1 u\left(\frac{X_{1,t}}{X_{7,t}}\right) + \gamma_2 u\left(\frac{X_{2,t}}{X_{7,t}}\right) + \gamma_3 u\left(\frac{X_{3,t}}{X_{7,t}}\right) + \gamma_4 u\left(\frac{X_{4,t}}{X_{7,t}}\right) + \gamma_5 u\left(\frac{X_{5,t}}{X_{7,t}}\right) + \gamma_6 u\left(\frac{X_{6,t}}{X_{7,t}}\right) + \gamma_7 u\left(X_{7,t}\right), \qquad (1)$$

with:

Y_t: Earned attention in period t;

$X_{1,t}$: Page impressions per different users in period t;

$X_{2,t}$: Quantity of commentary about product/brand in period t;

$X_{3,t}$: Brand mention in period t;

$X_{4,t}$: Total number of commentaries sent to friends in a short time (up to 5 minutes) after visiting the company website, in period t;

$X_{5,t}$: Total number of commentaries sent to friends in a medium time (from 5 minutes to a week) after visiting the company website, in period t;

$X_{6,t}$: Total number of commentaries sent to friends a long time (more than a week) after visiting the company website, in period t;

$X_{7,t}$: Total number of visitors to the company website in period t.

$\gamma_i, i = 1,\ldots,7$: weights of the different factors, defined according to the company preferences;

$u(\cdot)$: a utility function.

Page impressions (measured by $X_{1,t}$) denote that Internet users opened the company/brand webpage and also spent a minimum time looking at it (a minimum time of 30 seconds seems reasonable in order to consider that an impression is registered), so an increase in page impressions as a fraction of website visitors denotes an increase in the attention that the webpage was able to earn from the users. Commentaries left by the users ($X_{2,t}$) also denote an interest on the brand: if the user was willing to spend some time writing the commentary, this means that the brand earned enough attention from the user for her to be willing to spare that time. Brand mentions ($X_{3,t}$) mean that the user registered the brand in her memory and gave it enough relevance to refer it in a different context. Commentaries sent to friends ($X_{4,t}$, $X_{5,t}$, $X_{6,t}$) not only represent that the brand grabbed enough attention from the sender for her to send the comment to a friend but are also a measure of eWOM coming from a trusted source. We consider that commentaries sent a short time (up to 5 minutes), a medium time (more than 5 minutes and less than 1 week) and a long time (more than a week) after visiting the company website may reflect different levels earned attention, since they signal different levels of interest and different ways in which information is registered in the memory. All these variables are divided by the number of visitors to the company website $(X_{7,t})$, so they are included in the model as a fraction of the website visitors. Finally, the number of visitors to the company website $(X_{7,t})$ is also included as an independent variable, since it may also reflect in some way the attention earned by the company.

The values of these variables are converted into a utility by applying a function $u(\cdot)$. This function will be increasing, since we are assuming that the company wants to maximize the values of the variables or ratios that are considered. The rationale for using such a function is that the impact for the company of increasing the values of these variables will depend on their level – increasing page impressions from 1% to 2% of website visitors will usually have a larger marginal utility than increasing page impressions from 10% to 11% of website visitors. This can be introduced in the model by considering a concave function $u(\cdot)$ (that is, with $u''(\cdot) < 0$:). We can also consider different utility functions for different variables; however, in order to avoid cluttering the notation, in (1) we assume the same function is used for all variables, and we will follow this assumption hereafter. After converting the variables into utilities, the utilities of different variables are given a weight ($\gamma_i, i = 1,...,7$), defined according to the company preferences.

We assume that the company may influence the values of the variables $X_1 - X_7$ by undertaking marketing actions. There are n_A available actions and, in each period the company may undertake a subset of these actions. The chosen actions are constrained by a predefined budget.

Figure 1 illustrates the relations between the actions, variables and earned attention.

We assume that the company wants to maximize the average earned attention over n_t periods, $t = 1,...,n_t$. Each variable X_i will have an initial value $X_{i,0}$ and will change according to the actions undertaken by the company. There are n_A possible actions, $j = 1,...,n_A$ and in each period $t = 1,...,n_t$ each of these actions may be undertaken. Each action *j* has a cost C_j and its effect on the variables $X_1, X_2,...,X_7$ is defined incremental impacts $I_{j,1}, I_{j,2},...,I_{j,7}$.

Figure 1. Relation between actions, explanatory variables and earned attention

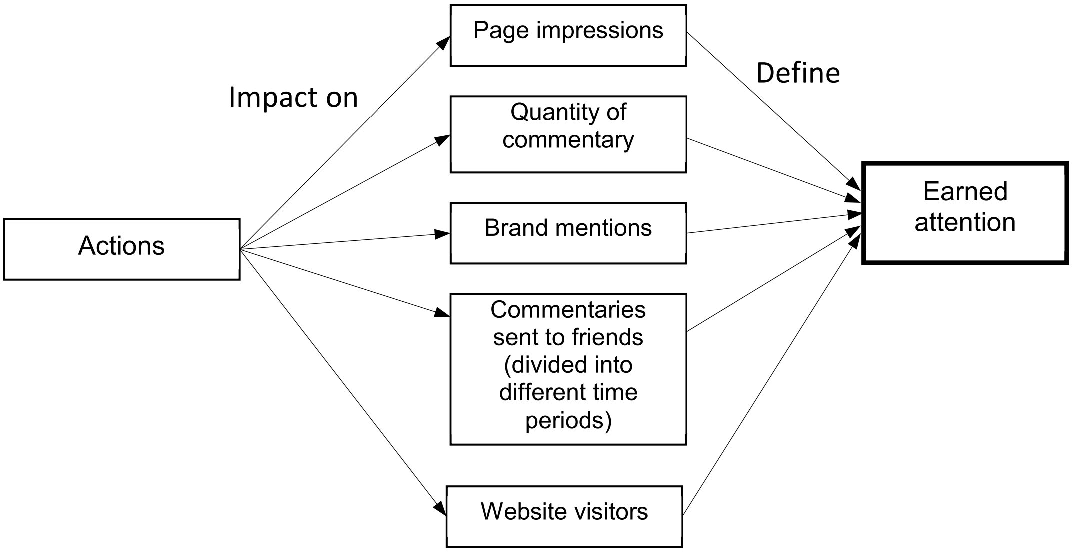

In each period the set of actions undertaken define a total impact $G_{i,t}$ on each variable $X_i, i = 1,...,7$. The choice of whether or not to undertake each action in each period is defined by the decision variables $y_{j,t}$:

$$y_{j,t} = \begin{cases} 0 \text{ if action } j \text{ is not undertaken in period } t \\ 1 \text{ if action } j \text{ is undertaken in period } t \end{cases} \tag{2}$$

Similarly to Naik et al. (1998), we assume that the use of the actions will cause repetition wearout, therefore reducing their impact in the subsequent periods, and interrupting the use of the actions will restore part of the impact. We use $d_{j,t} \in]0,1]$ to represent the effect of wearout on action *j* and we measure the impact of action *j* on variable X_i over period *t* as $I_{j,i} \cdot d_{j,t} \cdot y_{j,t}$. Notice that this means that the impact is $I_{j,i} \cdot d_{j,t}$ if action *j* is undertaken and zero if it is not. So, in period *t*, the total impact of the set of actions on variable X_i can be measured as

$$G_{i,t} = \sum_{j=1}^{n_A} I_{j,i} \cdot d_{j,t} \cdot y_{j,t} \tag{3}$$

In order to incorporate repetition wearout in the dynamics of $d_{j,t}$, we define that

$$d_{j,t+1} - d_{j,t} = -\omega_j \cdot y_{j,t} \cdot d_{j,t} + (1 - y_{j,t}) \cdot \delta_j \cdot (1 - d_{j,t}), \tag{4}$$

with ω_j representing the decay of the impact of action *j* due to wearout when this measure is used and δ_j representing the recovery of effectiveness of measure *j* when the action is not used in a given period.

The values of variables $X_1 - X_7$ in period *t* are influenced both by the previous values of these variables and by the global impact of the actions undertaken in period *t*. We write

$$X_{i,t} = f_i\left(X_{i,t-1}\right) + g_i\left(X_{i,t-1}, G_{i,t}\right), \tag{5}$$

where both functions $f_i\left(X_{i,t-1}\right)$ and $g_i\left(X_{i,t-1}, G_{i,t}\right)$ are increasing on their parameters.

In order to define a functional form for $f_i\left(X_{i,t-1}\right)$, we will consider that, in the absence of external impacts, there will be a constant decay of the values of variables $X_i, i = 1,...,7$, that is, a constant percentage reduction in the values of the variables. This means that the functions $f_i(\cdot), i = 1,...,7$ will consist on multiplying the previous value of the corresponding variable by a positive constant smaller than one. Therefore we can define

$$f_i\left(X_{i,t-1}\right) = \left(1 - k_i\right) X_{i,t-1}, i = 1,...,7 \tag{6}$$

with $k_i \in]0,1[, i = 1,...,7$ being a parameter representing the rate of natural decay of the values of variables $X_i, i = 1,...,7$, which must be estimated.

For defining $g_i(\cdot), i = 1,...,7,$ we assume that impact points have a bounded and monotonically increasing impact, and that initially the marginal impact is increasing and then it is decreasing. In order to achieve such a behaviour we define $g_i(\cdot)$ as Bass-type functions (Bass 1969):

$$g_i\left(X_{i,t-1}, G_{i,t}\right) = \alpha_i X_{i,t-1} \frac{1 - e^{-(p_i + q_i) G_{i,t}}}{1 + {}^{q}\!/_{p}\, e^{-(p_i + q_i) G_{i,t}}} \tag{7}.$$

In (7), $\alpha_i X_{i,t-1}$ is the supremum of the function $g_i(\cdot)$ and p_i and q_i are calibration parameters to be estimated.

Given these dynamics for the relevant variables, we intend to maximize the average earned attention in periods $t = 1,...,n_t$, which we define as:

$$\bar{Y} = \frac{1}{n_t}\sum_{t=1}^{n_t} Y_t \tag{8}$$

Finally, as we explained before, we assume that each action j has a cost C_j and that there is a global budget constraint B. So the choice of actions must take into account the following constraint.

$$\sum_{t=1}^{n_t}\sum_{j=1}^{n_A} C_j \cdot y_{j,t} \leq B \tag{9}$$

4. THE SOLUTION PROCEDURE

The function we aim to maximise is not continuous, since the choice of actions is defined by discrete variables, and it may possess several local optima, making it impossible to use derivative-based methods. Moreover, the solution space will usually be very large – in the examples shown in Section 5, we consider 12 actions available in either 10 or 15 periods, leading to solution spaces with about $6.2 \cdot 10^{10}$ and $1.5 \cdot 10^{16}$ solutions, respectively, making it impossible to use enumeration-based methods. Due to the non-linear impact of the variables on the objective function, we were unable to find an exact algorithm that could find the optimal solution in an acceptable time, so we resorted to a meta-heuristic.

Given that it was possible to take into account the characteristics of the problem in a local search procedure, we considered that the most promising meta-heuristics would combine a broad search of the solution space with the ability to perform local optimization. Therefore, we chose a Memetic Algorithm (MA), which combines the characteristics of genetic algorithms with the exploitation of the specific characteristics of this problem. In fact, Memetic Algorithms is the general designation of a family of meta-heuristics that try to mimic cultural evolution and are inherently concerned with exploiting all available knowledge about the problem under study. They are inspired by Neo-Darwinian's principles of natural evolution and Richard Dawkins' concept of a meme defined as a unit of cultural evolution that is capable of local refinements. Evolution is usually incorporated through a genetic algorithm-style procedure and local refinements through local search – Radcliffe and Surry define an MA as "a genetic algorithm incorporating local search" (Radcliffe and Surry 1994, p. 1). However, we must stress that the local search used in MAs often uses problem-specific information, making it more effective – this is the case with the application presented in this paper, which uses a local search procedure fitted to the problem at hand. MAs have been recognized as a powerful algorithmic paradigm for evolutionary computing and have proved to be successful for solving optimization problems in many contexts (e.g, Torn and Zilinskas 1989; Cowling et al. 2000; Ishibuchi et al. 2003).

4.1 General structure of the Memetic Algorithm

The MA uses a population of chromosomes, each one representing a strategy for maximizing the earned attention. The members of the population will also be referred to as "individuals". In a chromosome there will be a gene for each pair (action, period), indicating whether or not that action will be undertaken in that period: the gene will take the value one if the action is undertaken in that period and zero if it is not. So, each chromosome can be seen as a binary matrix with n_t lines and n_A columns. Each line will correspond to a time period and each column will represent an action available to the company. For example, if we have $n_t = 4$ periods and $n_A = 5$ actions, the chromosome corresponding to the strategy of undertaking actions 2 and 3 in period 1, action 4 in periods 2 and 3 and action 5 in period 4 can be represented by the following matrix:

$$\begin{array}{cc} & \text{Action} \\ & \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\ \text{Period} \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} & \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{array}$$

For computational convenience, chromosomes are not represented as matrices but as strings of genes. This means that the successive lines of the matrix are concatenated, leading to a string with $n_t \cdot n_A$ genes. In the previous example, the string would be

0 1 1 0 0 0 0 0 1 0 0 0 0 1 0 0 0 0 0 1

Each gene of the chromosome will therefore be one of the decision variables $y_{j,t}$ defined in (2), and for each chromosome we may use (1) and (3)-(8) to calculate the average earned attention that will be achieved through the corresponding strategy. The fitness of an individual will be this average earned attention, defined by (8), and is the measure of the quality of the corresponding solution.

The initial generation of individuals takes into account some information about the problem we are considering. We assume that the best individuals will represent solutions with a cost close to the available budget, and we try to generate individuals with these characteristics. In order to do that, we start by calculating the average cost of the available actions and we calculate the ratio between this cost and the available budget. When generating a new individual, the probability of each gene having the value one will be equal to this ratio, ensuring that the expected cost of each individual will be equal to the available budget.

In the MA, the population of individuals changes from a generation to the next by applying several operators to the individuals. The first operator to be applied is the selection operator, which randomly chooses two individuals to become parents of two new individuals of the next population. We use roulette wheel selection, meaning that the probability of an individual being selected is proportional to is fitness.

After the parents are chosen, the MA randomly determines if the new individuals are equal to the parents or if it will result from a combination of the parents. The probability of the new individuals being a combination of the two parents is termed the crossover probability, and the combination of the two parents is performed through the use of the crossover operator. We use 1-point crossover, meaning that the crossover operator randomly chooses a locus and exchanges the subsequences before and after that locus between the parent chromosomes to create the new individuals.

In order to introduce some innovations in the population, a mutation operator is used. This operator considers all genes of each chromosome and, with a probability termed the "mutation probability", it changes their values. Finally, we introduce elitism in the algorithm by reintroducing, in the population, the best individual of the previous generation, whenever the best element of the current generation is worse than that one.

One problem with the application of an MA to this problem has to do with the fact that some chromosomes will correspond to non-admissible solutions, that is, solutions in which the budget constraint does not hold. We handle this problem through the local search component of the MA. The local search starts by analyzing the chromosome in order to check whether the corresponding solution is feasible. If the solution is not feasible, the local search consecutively removes actions that are being undertaken (that is, changes genes with the value of one to zero), until the budget constraint holds. Afterwards, further local optimization is performed. We describe the local search procedure in the next section.

The introduction of local search in an evolutionary algorithm may sometimes lead to premature convergence to local optima that are not global optima. In order to avoid this problem, we follow a strategy based on the adjustment of the mutation rate. We define two mutation rates: a base rate and an increased rate. The base rate is used by default. When there are indications that the diversity in the population is becoming too low, we switch to the increased rate. This increased rate ensures faster changes in the chromosomes of the population, leading to an increase in the diversity. When enough diversity is restored, the base mutation rate is used again.

The variance of the fitness values is used as an indicator of the diversity of the population. This is a very rough indicator, since two individuals with similar fitness may have very different genetic structure. However, it is an indicator that can be computed quickly and that is sure to identify extreme cases of a population composed by very similar individuals. Other alternatives based on the direct calculation of a "similarity coefficient" for each pair of individuals were left out due to the computational burden that they would entail.

In each generation, the MA checks whether the base mutation rate or the increased mutation rate should be used. In order to do that, the variance of the fitness values of the individuals in the population is calculated, and it is then compared with the maximum variance hitherto achieved. If the current variance of the fitness values is smaller than a pre-defined percentage of the maximum variance previously achieved, it is considered that the diversity declined enough to justify the use of the increased mutation rate. Otherwise, the base mutation rate is used.

These procedures – selection, crossover, mutation, local search and elitism – are repeated for evolving the population and defining the successive generations. The stopping criterion we use is based on the number of generations: we define a number of generations and stop the algorithm when this number is reached. The solution returned by the MA will correspond to the best individual of the last generation.

4.2 The local search procedure

We incorporate in the MA a specific local search procedure that takes into account the characteristics of the problem. The local search procedure resorts to a ranking of the genes based on the ratio between the "attractiveness" of the gene and the cost of the corresponding action. In order to measure the "attractiveness" of a gene we consider a base situation in which all genes have a zero value (no actions are undertaken) and we calculate the impact on the fitness value of changing the gene value to one. Somewhat more formally, let us define that $a_{j,t}$ is the difference between the fitness of the chromosome in which the only action undertaken is action *j* in period *t* and that of the chromosome corresponding to the case in which no actions are undertaken. Then the "attractiveness ratio" of the gene corresponding to undertaking action *j* in period *t* is defined as:

$$r_{j,t} = a_{j,t} / C_j \,, \tag{10}$$

with C_j being the cost of action *j*, as defined in Section 3. Using the matrix structure of the chromosomes introduced in the previous subsection, and considering a very small example with just three actions and two periods, $a_{2,1}$ is the difference between the fitness of chromosome $\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and that of $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$. $r_{2,1}$ is then defined as the result of $a_{2,1} / C_2$

We consider that genes with a larger $r_{j,t}$ are more promising since they are expected to lead to a greater increase in the fitness function for each unit of cost. There are four components in this procedure, which we will now describe. We will illustrate them based on an example with three actions and two periods, in which we assume $r_{1,1} = 100; r_{2,1} = 90; r_{3,1} = 95; r_{1,2} = 80; r_{2,2} = 70; r_{3,2} = 75$, $C_1 = 10; C_2 = 8; C_3 = 7$ and a budget constraint $B = 27$. The components of the procedure are executed in the following order:

1. If the individual exceeds the available budget, successively remove actions (that is, change genes with the value of one to zero) until the budget constraint holds, considering the genes in increasing order of $r_{j,t}$. As an example, assume that the current solution is represented by chromosome $\begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix}$. The total cost is

$2 \cdot C_1 + C_2 + 2 \cdot C_3 = 42$, exceeding the available budget. The smallest $r_{j,t}$ for a gene with a value of one is $r_{3,2} = 75$, so we change the value of that gene to zero. The solution corresponding to the chromosome still has a cost of 35, so we find the smallest $r_{j,t}$ for genes with a value of one, which is $r_{1,2} = 80$. We end up with the chromosome $\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$, which has a cost of 25, within the available budget.

2. For each gene with value one, check whether changing it to zero would increase or decrease the fitness of the individual. If it would increase its fitness, make the corresponding change to the individual. The checks are made in increasing order of $r_{j,t}$. As an example, if the current solution corresponds to chromosome $\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$, we would consecutively check whether $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$, $\begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ or $\begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$ would have a higher fitness, and change the chromosome if we would find one with better fitness.
3. For each gene with value zero (an action that is not being undertaken in a given period), check whether the corresponding action can be undertaken without exceeding the budget. If it can, and if that leads to an increase in the fitness of the individual, then change the gene value to one. The genes are considered in decreasing order of $r_{j,t}$. For example, if the current solution corresponds to chromosome $\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$, then the largest $r_{j,t}$ is $r_{3,1} = 95$, and changing the corresponding gene to one does not violate the budget constraint. We would compare the fitness of the previous chromosome with that of $\begin{bmatrix} 1 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$ and, if the latter would have a higher fitness, we would replace the previous one by it.

4. For each gene with value one (an action being undertaken in a given period), try to change the period in which the action is being undertaken. If it is possible to improve the fitness of the individual in this way, then make the corresponding changes. As an example, consider the chromosome $\begin{bmatrix} 1 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$, in which action 1 is undertaken in periods 1 and 2, and action 3 is undertaken in period 1. Since action 1 is undertaken in all periods, we cannot try to use it in different periods. Action 2 is never used, so there cannot be a change in the period in which it is undertaken. However, we can change the period in which action 3 is undertaken. For that, we would compare the fitness of the previous chromosome with that of $\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix}$ and, if the latter would have a better fitness, we would replace the previous one by it.

5. COMPUTATIONAL ANALYSIS USING SIMULATED DATA

In order to get some indications about the behaviour of the model and the solution procedure, the model and the solution procedure were implemented in Pascal programming language (using the Free Pascal compiler) and a computational analysis was performed based on simulated data. In the tests, we considered isoelastic utility functions with the form

$$u(x) = \begin{cases} \dfrac{x^{1-\eta}-1}{1-\eta}, \text{ if } \eta \neq 1 \\ \ln(x), \text{ if } \eta = 1 \end{cases} \tag{11}$$

where x is the value of the measure whose utility is being calculated and η is a parameter that can be interpreted as the degree of risk aversion or the speed of the decrease in the marginal utility of x. By looking at (1), it is possible to realize that the arguments of $u(\cdot)$ can take any non-negative value. In order to have utilities starting in zero – that is, in order to have $u(0) = 0$ – we add 1 to the value of each measure before calculating the utilities. Additionally, since the last term of (1) is the only one that does not apply the utility to a quotient, in that term we divide $X_{7,t}$ by 100 000, in order to avoid having values of very different magnitudes.

We defined 8 types of problems, based on three characteristics: rate of natural decay of the values of variables $X_i, i=1,...,7$ ($k_i = 0.75, i=1,...,7$ - large decay and $k_i = 0.5, i=1,...,7$ - small decay), number of periods (10 periods and 15 periods) and parameter η ($\eta = 0$ - constant marginal utility and $\eta = 2$ - decreasing marginal utility). For each type of problem, five problems were generated randomly. Each problem was then solved eight times using the MA. In order to allow us to assess whether it is worthwhile to use an MA for this problem, a standard local search procedure was used for comparison purposes (assigning the same execution time to each method, in order to allow a fair comparison of results). The results were then compared.

Based on some preliminary computational experiments, we used the following parameters in the MA

- Population size: 80
- Number of generations: 400
- Crossover rate: 40%
- Base mutation rate: 2.5%
- Increased mutation rate: 20%
- Percentage of maximum variance previously achieved that triggers the use of increased mutation rate: 20%

5.1 Data used in the analysis

In order to define the problems, several parameters were given values we considered reasonable, and kept constant in all problems. This was the case of the parameters shown in Table 1. For the values of p_i and q_i (used in expression (7)) it was considered that they should be calibrated in such a way that 5 impact points would always lead to $g_i\left(X_{i,t-1}, 5\right) = 25\%\alpha_i X_{i,t-1}$

and 10 impact points would lead to $g_i\left(X_{i,t-1},10\right)=50\%\alpha_i X_{i,t-1}$. This led to $p_i=0.04621$ and $q_i=0.09242$.

Table 1 – Values of the parameters kept constant in all problems.

Parameter	Value	Parameter	Value
$X_{1,0}$	10 000	$p_i, i=1,...,7$	0.04621
$X_{2,0}$	1500	$q_i, i=1,...,7$	0.09242
$X_{3,0}$	1500	$\omega_i, i=1,...,7$	0.4
$X_{4,0}$	1000	$\delta_i, i=1,...,7$	0.1
$X_{5,0}$	500	$\alpha_i, i=1,...,7$	150%
$X_{6,0}$	500		
$X_{7,0}$	25 000		

The weights of the different factors, $\gamma_i, i=1,...,7$, should be defined according to the company preferences. In the analysis, we allowed these weights to change randomly by defining reasonable ranges of values and generating each weight in each problem as a draw from a uniform distribution. This way, the range of results achieved in the analysis will reflect the ranges defined for these weights. These ranges are shown in Table 2.

Table 2 – Ranges of the weights used in the problems.

Weight	Range of values	Weight	Range of values
γ_1	[0.75, 1.25]	γ_5	[20.0, 40.0]
γ_2	[7.5, 12.5]	γ_6	[20.0, 40.0]
γ_3	[7.5, 12.5]	γ_7	[0.35, 0.65]
γ_4	[12.0, 18.0]		

The rate of natural decay of the values of variables $X_i, i=1,...,7$, the number of periods and the parameter η of the utility functions define the type of each problem. For the number of periods, the values $n_t=10$ and $n_t=15$ were used. For the rate of decay of the values of variables $X_i, i=1,...,7$, we considered large decay for all variables, $k_i=0.75, i=1,...,7$, and small decay for all variables, $k_i = 0.5, i = 1,...,7$. For the parameter η, we used $\eta = 0$ for constant marginal utility and $\eta = 2$ for decreasing marginal utility.

In all problems we considered that there were 12 actions available. We defined three different types of actions, and in each problem there are four available actions of each type. Each type of action is defined by a range of costs and a range of values of $I_{j,i}$, representing the impact of an action *j* on each variable $X_i, i = 1,...,7$. The characteristics of each action used in each problem are defined as random draws from these ranges of values. The types of actions are denoted as A, B and C, and the ranges used for each of these parameters are shown in Table 3.

Table 3 – Ranges of the parameters' values for each type of action.

Type of action	Cost (C_j)	$I_{j,1}$	$I_{j,2}$	$I_{j,3}$	$I_{j,4}$	$I_{j,5}$	$I_{j,6}$	$I_{j,7}$
A	[200, 300]	[4.0, 5.0]	[1.0, 2.0]	[1.0, 2.0]	[1.0, 2.0]	[1.0, 2.0]	[1.0, 2.0]	[1.0, 1.5]
B	[200, 300]	[1.0, 2.0]	[1.0, 2.0]	[3.0, 4.0]	[1.0, 2.0]	[1.0, 2.0]	[1.0, 2.0]	[1.0, 1.5]
C	[350, 450]	[1.0, 2.0]	[2.0, 3.0]	[1.0, 2.0]	[2.0, 3.0]	[2.0, 3.0]	[2.0, 3.0]	[1.0, 1.5]

By looking at Table 3 we can see that actions of type A tend to have a larger impact on page impressions, actions of type B have a larger impact on brand mentions and actions of type C have larger impact on commentary production and dissemination. Given that actions of type C have a broader impact, we decided to assign a larger cost to these actions.

Finally, in order to define the available budget, we considered 600 monetary units per period. This means that the budget depends on the type of problem. The budget is 6000 for problems with 10 periods and 9000 for problems with 15 periods.

5.2 Memetic algorithm comparator

In order to assess the performance of the MA, a standard Steepest Ascent (SA) heuristic was applied to the generated problems, and the results were compared to those of the MA. The

SA procedure also uses the same solution representation as the MA, that is, a solution is a binary string with $n_t \cdot n_A$ elements. The heuristic starts by considering an initial solution generated randomly. Afterwards the heuristic tries to change the state of all the genes in the chromosome. It starts with the first gene and changes its value (from 1 to 0 or from 0 to 1, depending on the initial value). The fitness of the new solution is compared with the fitness of the previous one and the current solution becomes the one with the highest fitness value. This procedure is repeated for all genes, until the last one is reached. If the initial solution was changed during this pass, then a new pass through all the genes is performed, and the algorithm goes on until a complete pass is performed without improving the fitness value.

In order to allow the SA method to avoid unfeasible solutions, the evaluation of solutions for this method penalizes such solutions: each unfeasible solution suffers a constant penalization of 1000 units of the fitness function, and is also penalized in 1000 additional units per unit of excessive cost. The constant component of the penalization prevents the procedure from achieving solutions whose cost exceeds the budget by a small margin, and the variable component leads the procedure to successively correct excessive costs. Additionally, in order to allow a fair comparison of the results of the SA to those of the MA, the procedure is successively repeated, with different initial solutions, for the same time it took to execute the MA.

5.3 Results and analysis

Table 4 reports the average fitness achieved for each problem type and the standard deviation of the fitness values obtained in different runs, as well as the differences in average fitness for the two methods. Since the fitness values of different problems have different magnitudes, we have chosen to compare the average percentage differences in fitness for each pair of methods. So, for each random problem we calculate the average difference of fitness over the eight runs for each pair of methods, and then, for each type of problem, we calculate the average difference for the generated problems of that type. For the standard deviation, we start by calculating the standard deviation over the 8 runs of each problem. We then divide it by the average fitness and average the values thus obtained over the problems of each type.

Table 4 – Average fitness and standard deviation of fitness for the different types of problems and differences of performance for the two methods.

Problem type	Average fitness		Standard deviation		Difference
	MA	SA	MA	SA	MA over SA
Small decay, 10 periods, $\eta = 0$	9.61	8.60	0.73%	0.96%	10.55%
Small decay, 10 periods, $\eta = 2$	7.17	6.74	0.20%	0.70%	5.93%
Small decay, 15 periods, $\eta = 0$	16.29	12.98	1.55%	1.70%	19.45%
Small decay, 15 periods, $\eta = 2$	9.44	8.57	0.26%	0.67%	9.19%
Large decay, 10 periods, $\eta = 0$	33.71	21.89	0.34%	5.88%	35.07%
Large decay, 10 periods, $\eta = 2$	10.81	9.87	0.29%	1.04%	8.70%
Large decay, 15 periods, $\eta = 0$	133.32	55.73	0.44%	6.67%	58.18%
Large decay, 15 periods, $\eta = 2$	16.31	14.25	0.34%	1.37%	12.68%

Average fitness calculated over all the problems of each type; standard deviation calculated for each problem, expressed as a percentage of the average fitness for the problem and averaged for each problem type. Differences calculated for each problem, expressed as a percentage and averaged for each problem type.

The MA clearly outperforms the SA method for all types of problems, the differences being much larger for the problems with 15 periods. This was to be expected, since those problems have a larger search space, and the local search methods will have more difficulty in finding good solutions, while the MA will use its ability to explore promising regions of the search space to achieve solutions with large values of the fitness function. It is also clear from Table 4 that the differences in performance are larger when we consider constant marginal utility ($\eta = 0$). This is due to the fact that considering decreasing marginal utility tends to reduce the performance differences between strategies in general, therefore reducing also the magnitude of improvements provided by better solution methods.

In Table 4, we can see that the standard deviations obtained for the MA method are quite small (usually smaller than 1%). This provides an indication that all the runs are reaching solutions that lead to a similar fitness. We can also see that the average fitness is larger for problems with a larger number of periods (15 periods), larger decay and constant marginal utility ($\eta = 0$). Problems with a larger number of periods provide more opportunities for building strategies that lead to a larger value of the average earned attention, particularly since the budget is bigger for a higher number of periods, therefore leading to a higher fitness. A larger decay allows a faster decline in the denominator of the quotients in the first terms of (1), allowing the algorithms to define strategies that increase the values of these quotients more significantly, and leading to larger fitness. Finally, the functions $u(\cdot)$ used for decreasing marginal utility lead to smaller utility values, therefore constant marginal utility leads to higher fitness.

Tables 5 and 6 show some characteristics of the best solutions achieve by the MA. In Table 5 we can see that the type of actions that lead to better solutions depend a lot on the assumption of constant or decreasing marginal utility. With constant marginal utility, actions of type A are clearly favoured; when there is decreasing marginal utility, actions of type C are the most used ones, and actions of type A the ones that are used less often. In fact, actions of type A are the ones that allow a larger impact over a single variable: the number of page impressions. With constant marginal utility the best solutions are the ones that build very large values of this variable, by using actions of type A multiple times, since the influence of the variable in the fitness function does not decrease with its level. On the other hand, with decreasing marginal utility, increasingly larger values of a variable have a diminishing influence in the fitness function. So, the best solutions tend to be the more balanced ones. Actions of type C, with a medium impact over several variables, tend to be the ones that are used most often. We can conclude that the characteristics of the utility function have a very significant impact on the type of strategies that perform best. In general, the more significant is the decrease in marginal utility with the level of the variable, the greater the preference for a portfolio of actions that leads to a balanced impact over all the variables that are considered.

Table 5 – Average number of times that an action of each type is used, in the best solutions found by the MA for each type of problem

		10 periods		
Type of action	Small decay, $\eta = 0$	Small decay, $\eta = 2$	Large decay, $\eta = 0$	Large decay, $\eta = 2$
Type A	3.66	0.89	5.58	0.68
Type B	0.96	1.80	0.02	1.53
Type C	0.86	2.16	0.00	2.45
		15 periods		
Type of action	Small decay, $\eta = 0$	Small decay, $\eta = 2$	Large decay, $\eta = 0$	Large decay, $\eta = 2$
Type A	5.58	1.43	7.76	1.06
Type B	1.61	2.78	0.24	2.27
Type C	1.11	3.12	0.16	3.68

Table 6 – Average number of actions used in each period, in the best solutions found by the MA for each type of problem.

		10 periods		
Period	Small decay, $\eta = 0$	Small decay, $\eta = 2$	Large decay, $\eta = 0$	Large decay, $\eta = 2$
1	2.35	2.78	1.79	2.05
2	2.43	2.93	1.72	2.08
3	2.38	2.83	1.38	2.15
4	2.70	2.90	1.64	2.65
5	2.63	2.73	1.95	2.48
6	2.83	2.55	1.97	2.35
7	2.75	1.98	2.13	2.13
8	2.38	0.73	2.87	1.80
9	1.35	0.00	3.51	0.93
10	0.13	0.00	3.41	0.00
		15 periods		
Type of action	Small decay, $\eta = 0$	Small decay, $\eta = 2$	Large decay, $\eta = 0$	Large decay, $\eta = 2$
1	2.10	2.65	1.25	2.03
2	2.08	2.85	1.58	2.15
3	2.15	2.78	1.20	2.20
4	2.30	2.98	1.50	2.48
5	2.15	2.93	1.60	2.43
6	2.40	2.75	1.70	2.53
7	2.68	2.70	1.60	2.43
8	2.68	2.68	1.93	2.68
9	2.75	2.40	2.10	2.40
10	2.78	2.08	2.03	2.38
11	2.75	1.75	2.13	2.08
12	2.28	0.75	2.40	1.65
13	2.08	0.03	3.00	0.63
14	1.58	0.00	3.90	0.00
15	0.45	0.00	4.70	0.00

In Table 6 we can see that the average number of actions undertaken in each period tends to increase slightly in the beginning and decrease in the last periods. This latter fact was to be expected, since actions undertaken in the last periods will have an effect for a limited time, and thus they may contribute less to the average earned attention. There are, however, two exceptions to this behaviour: the cases in which there is constant marginal utility and large decay in the effectiveness of an action due to its repeated use. As explained before, in case of constant marginal utility, the best strategies tend to favour a focus on using actions that have a large impact over a variable. If, additionally, there is a large decay in the values of the variables, this allows the denominator of the quotient present in the first terms of (1) to reach small values in the latest periods. Increasing the denominator in such periods leads to large values of the quotient, so the best strategies try to maximize the values of the variable in the numerator in the latest periods, repeatedly using the actions that lead to its increase.

Finally, let us show an example of the best solution found by the MA in a run of a problem with 10 periods, decreasing marginal utility and large decay:

$$\begin{array}{ccc}
 & & \text{Action} \\
 & & \begin{array}{cccccccccccc} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 \end{array} \\
\text{Period} & \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \\ 10 \end{array} & \begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}
\end{array}$$

This example allows us to see some characteristics common to most of the best solutions found for these problems. First, we can notice that some actions tend to be favoured and used several times, while others tend to be disregarded. A second relevant fact is that, due to the existence of repetition wearout, the best solutions tend to avoid using the same actions in multiple consecutive periods. In this example, no action is used in two consecutive periods; in other cases, an action is used on two consecutive periods, but rarely on more than that.

6. CONCLUSIONS

In this research study, we have tried to test and validate a model that attempts to maximise the amount of consumer earned attention gained by a brand. The explanatory variables comprising the research model were grouped into online traffic measures, conversation flows amongst member of a particular social network, subjected to three different time frames, and two associated measures related consumer-brand relationships. We have used a Memetic Algorithm (MA) to test the assumptions included in the model, which itself contains characteristics of genetic algorithms coupled with the exploration of the specific facets of the research study problem. A number of key features were embedded in the algorithm, such as the use of a common roulette selection procedure-1-point crossover and mutation operators in the MA, as well as using the variance of fitness values as an indicator of the diversity of the population. In order to assess the performance of the MA, an alternative local search method was utilized – the Steepest Ascent (SA) technique. The MA always outperforms the SA heuristic.

The major findings stemming from the validation of the research model refer to the fact that the type of utility function that is considered may have a significant impact on the characteristics of the best strategies. Under constant marginal utility, actions focused on significantly increasing a single variable will be preferred. If the marginal utility is decreasing, then the choice of actions of different types will be more balanced, with actions with a broader impact being used more often. It was also found that when the repetitive wearout effect may have a significant impact, the intermittent use of marketing actions may be advised, instead of using these actions for long consecutive periods. Implications for practitioners are relevant on how to leverage on an effective social media strategy in order to increase earned attention and influence customer behaviour.

REFERENCES

Ahuja, V., & Medury, Y. (2010). Corporate blogs as e-CRM tools – Building consumer engagement through content management. *Journal of Database Marketing & Customer Strategy Management*, 17(2), 91–105.

Alba, J., & Hutchinson, W. (1987). Dimensions of Consumer Expertise. *Journal of Consumer Research,* 13 (4), 411–454.

Anderson, E., & Mittal, V. (2000). Strengthening the Service Profit Chain. *Journal of Service Research, 3*(2), pp. 107-120

Barnatt, C. (2001). The second digital revolution. *Journal of General Management, 27*(2), pp. 1-16.

Bass, F.M. (1969). A New Product Growth for Model Consumer Variables. *Management Science*, 15(5), 215-227.

Bergkvist, L. & Bech-Larsen, T. (2010). Two studies of consequences and actionable antecedents of brand love. *Journal of Brand Management*, 17, 504-518.

Berthon, P.R., Pitt, L.F., McCarthy, I.P., & Kates, S.M. (2007). When customers get clever: Managerial approaches to dealing with creative consumers. *Business Horizons*, 50(1), 39-47.

Bolotaeva, V., & Cata, T. (2011). Marketing opportunities with social networks. *Journal of Internet Social Networking and Virtual Communities*. http://www.ibimapublishing.com/journals/JISNVC/2011/409860/409860.pdf. Accessed 29 July 2014.

Bowden, J. L. H. (2009). The process of customer engagement: A conceptual framework. *Journal of Marketing Theory and Practice, 17*(1), 63–74.

Brodie, R.J., Hollebeek, L. D., Juric, B., & Llic, A. (2011a). Customer Engagement: Conceptual Domain, Fundamental Propositions, and Implications for Research. *Journal of Service Research, 14*(3), 252-271.

Cowling, P., Kendall, G., & Soubeiga, E. (2000). A hyperheuristic approach to scheduling a sales summit. In E. Burke, & W. E. Editors (Eds.), *Proc. Theory of Automated Timetabling PATAT 2000, Lecture Notes in Computer Science* (pp. 176–190). Berlin, Germany: Springer-Verlag.

Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience.* New York: Harper and Row.

Dennison, G., Bourdage-Braun, S., & Chetuparambil, M. (2009). Social Commerce Defined. IBM White Paper, #23747, November.

Engel, J.F., Blackwell, R.D., & Kegerreis, R.J. (1969). How information is used to adopt an innovation. *Journal of Advertising Research*, 9(4), 3-8.

Feick, L.F., & Price, L.L. (1987). The market maven: a diffuser of marketplace information. *Journal of Marketing*, 51(1), 83–97.

Gambetti, R. C., & Graffigna, G. (2010). The concept of engagement: A systematic analysis of the ongoing marketing debate. *International Journal of Market Research, 52*(6), 801-826.

Gilbert, S.J. (2009). Social network marketing: What works? Sales&MktgWatch, 9(7). http://hbswk.hbs.edu/item/6187.html. Accessed 29 July 2014.

Gilly, M.C., Graham, J.L., Wolfinbarger, M.F., & Yale, L.J. (1998). A dyadic study of interpersonal information search. *Journal of the Academy of Marketing Science*, 26(2), 83-100.

Godes, D., & Mayzlin, D. (2004). Using online conversations to study word-of-mouth communication. *Marketing Science*, 23(4), 545-560.

Godinho, P., Moutinho, L., & Silva, M. (2015) A model for optimising earned attention in social media, in Moutinho, L. and Huarng, K.-H. (Eds.) *Quantitative Modelling in Marketing and Management, 2nd edition* (pp. 423-456). World Scientific.

Godwin, M. (1994). Nine principles for making virtual communities work. *Wired 2.06*, 72–73.

Gummerus, J., Liljander, V., Weman, E., & Pihlström, M. (2012). Customer engagement in a Facebook brand community. *Management Research Review*, 35(9), 857–877.

Handelsman, M. M, Briggs, W. L., Sullivan, N., & Towler, A. (2005). A measure of college student course engagement. *The Journal of Educational Research, 98*(3), 184-191.

Hanna, R., Rohm, A., & Crittenden, V.L. (2011). We're all connected: The power of the social media ecosystem. *Business Horizons*, 54(3), 265-273.

Heath, R. (2009). Emotional engagement: how television builds big brands at low attention. *Journal of Advertising Research, 49*(1), pp. 62-73.

Hennig-Thurau, T., Gwinner, K.P., Walsh, G., & Gremler, D.D. (2004). Electronic word-of-mouth via consumer-opinion platforms: what motivates consumers to articulate themselves on the internet? *Journal of Interactive Marketing*, 18(1), 38–52.

Heskett, J., Jones, T., Loveman, G., Sasser, E., & Schlesinger, L. (1994). Putting the Service Profit Chain to Work. *Harvard Business Review, 72*(2), 164-171.

Holbrook, M.B. and Hirschman, E.C. (1982). The Experiential Aspects of Consumption: Consumer Fantasy, Feelings and Fun, *Journal of Consumer Research*, 9(2), 132-140.

Hollebeek, L. D. (2011a). Demystifying customer brand engagement: Exploring the loyalty nexus. *Journal of Marketing Management, 27*(7-8), 785-807

Hung, K.H., & Li, S.Y. (2007). The influence of eWOM on virtual consumer communities: social capital, consumer learning, and behavioral outcomes. *Journal of Advertising Research*, 47(4), 485–495.

Ishibuchi, H., Yoshida, T., & Murata, T. (2003). Balance between genetic search and local search in memetic algorithms for multiobjective permutation flowshop scheduling. *IEEE Transactions on Evolutionary Computation*, 7(2), 204–223.

Javornik, A., & Mandelli, A. (2012). Behavioral perspectives of customer engagement: An exploratory study of customer engagement with three swiss FMCG brands. *Journal of Database Marketing & Customer Strategy Management, 19*(4), 300-310.

Johnson, C., & Mathews, B. (1997), The Influence of Experience on Service Expectations, *International Journal of Service Industry Management,* 8 (4), 290–305.

Kastelein, R., & Rempt, D. (2010) Social TV the virtual water cooler? Association of International Broadcaster's. *The Channel Magazine*, (2), 30-31. Retrieved from http://www.appmarket.tv/documents/cat_view/162-corporate-research.html

Katz, E., & Lazarsfeld, P.E. (1955). Personal Influence: The Part Played by People in the Flow of Mass Communications. New York: Free Press.

Kim, A. (2000). *Community Building on the Web.* Peachpit Press, Berkeley, CA.

Koh, J., Kim, Y.G., Butler, B., & Bock, G.W. (2007). Encouraging Participation in Virtual Communities. *Communications of the ACM*, 50 (2), 69-73.

Marsden, P.V., & Podolny, J. (1990). Dynamic Analysis of Network Diffusion Processes, Weeise, J. and Flap, H. (Eds.) *Social Networks through Time*, Uthrecht, The Netherlands: ISOR/Rijksuniversiteit Utrecht.

Matilla, A., & Wirtz, J. (2002). The Impact of Knowledge Types on the Consumer Search Process: An Investigation in the Context of Credence Services, *International Journal of Service Industry Management,* 13 (3), 214–230.

Mersey, R. D., Malthouse, E. C., & Calder, B. J. (2010). Engagement with online media. *Journal of Media Business Studies, 7*(2), pp. 39-56

Moreau, P., Lehmann, D., & Markman, A. (2001). Entrenched Knowledge Structures and Consumer Response to New Products, *Journal of Marketing Research,* 38 (1), 14–30.

Naik, P.A., Mantrala, M.K., & Sawyer, A.G. (1998). Planning media schedules in the presence of dynamic advertising quality. *Marketing Science*, 17(3), 214-235.

Neff, J. (2007). OMD proves the power of engagement. *Advertising Age, 78*(27), pp.3-4

O'Brien, H. L., & Toms, E. G. (2010). The Development and Evaluation of a Survey to Measure User Engagement. *Journal of the American Society for Information Science and Technology, 61*(1), pp. 50–69.

Ojiako, U., Chipulu, M., & Graesser, A. (2012). Correlating service touch-point preferences with engagement parameters. *Industrial Management + Data Systems, 112*(5), 766-785.

Pagani, M., & Mirabello, A. (2011). The Influence of Personal and Social-Interactive Engagement in Social TV Web Sites. *International Journal of Electronic Commerce, 16*(2), pp. 41-68.

Patterson, P. (2000). A Contingency Approach to Modeling Satisfaction with Management Consulting Services, *Journal of Service Research,* 3 (2), 138–153.

Porter, C. E., Donthu, N., MacElroy, W. H., & Wydra, D. (2011). How to Foster and Sustain Engagement in Virtual Communities. *California Management Review, 53*(4), pp. 80-110

Prahalad, C.K., & Ramaswamy, V. (2004). Co-creation experiences: The next practice in value creation, *Journal of Interactive Marketing*, Vol. 18, 3, 5-14.

Radcliffe, N.J., & Surry, P.D. (1994). Formal Memetic Algorithms. In T.C. Fogarty (Ed.), *Evolutionary Computing, Lecture Notes in Computer Science 865* (pp. 1-16). Berlin, Germany: Springer-Verlag.

Reichheld, F. (2003). The One Number You Need to Grow. *Harvard Business Review, 81*(12), pp. 46–54.

Reichheld, F., & Sasser, W. E. (1990). Zero defects: Quality comes to Services. *Harvard Business Review*(Sept-Oct), pp. 105–111.

Sashi, C. M. (2012). Customer engagement, buyer-seller relationships, and social media. *Management Decision, 50*(2), pp. 253 – 272.

Soderlund, M. (2002). Customer Familiarity and Its Effects on Satisfaction and Behavioral Intentions, *Psychology & Marketing,* 19(10), 861–880.

Sonnier, G.P., McAlister, L., & Rutz, O.J. (2011). A dynamic model of the effect of online communications on firm sales. *Marketing Science*, 30(4), 702-716.

Torn, A., & Zilinskas, A. (1989). Global Optimization, Vol. 350, Lecture Notes in Computer Science. Berlin, Germany: Springer-Verlag.

Tripathi M.N. Vilakshan (2014). Customer Satisfaction and Engagement: Customer retention strategies for brand managers. *The XIMB Journal of Management,* Vol. 11 (1), p123-134.

Van Doorn, J., Lemon, K. N., Mittal, V., Nass, S., Pick, D., Pirner, P., & Verhoef, P. C. (2010). Customer Engagement Behavior: Theoretical Foundations and Research Directions. *Journal of Service Research*, 13(3), 253–266.

Vollmer, C., & Precourt, G. (2008). Always On: Advertising, Marketing, and Media in An Era of Consumer Control. New York: McGraw-Hill.

Weston, R. (2008). 7 Social Networking Strategies. Entrepreneur.com. http://www.entrepreneur.com/technology/bmighty/article191312.html. Accessed 29 July 2014.

Williams, R.L., & Cothrel, J. (2000). Four Smart Ways to Run Online Communities. *Sloan Management Review*, 41 (4), 81-91.