# Building Agent-Based Decision Support Systems for Word-Of-Mouth Programs. A Freemium Application

Manuel Chica[*,a], William Rand[b]

[a]Department of Computer Science - IN3, Open University of Catalonia, 08018 Barcelona, Spain
[b]Robert H. Smith School of Business, University of Maryland, USA

## Abstract

Marketers have to make decisions on how to implement word-of-mouth (WOM) programs and a well-developed decision support system (DSS) can provide them with valuable assistance. The authors propose an agent-based framework that aggregates social network-level individual interactions to guide the construction of a successful DSS for WOM. The framework presents a set of guidelines and recommendations to involve stakeholders, follow a data-driven iterative modeling approach, increase validity through automated calibration, and understand the DSS behavior. This framework is applied to build a DSS for a freemium app, where premium users discuss the product with their social network and promote the viral adoption. After its validation, the agent-based DSS forecasts the aggregate number of premium sales over time and the most likely users to become premium in a near future. The experiments show how the DSS can help managers by forecasting premium conversions and increasing the number of premiums via targeting and rewarding policies.

*Key words:* Word-of-mouth, Marketing Decision Support Systems, Agent-based Modeling, Targeting and Referrals, Freemium Business Model

One day a brand manager walks into her office, and comes up with what she thinks is a great idea, she should send out a designer, one-of-a-kind, unique t-shirt to her most valuable customers. She hopes that customers who receive these t-shirts will talk about them with their friends, and those friends will be more likely to become customers themselves. She knows that word-of-mouth (WOM) can be a powerful force for marketing (Trusov et al. 2009), and she hopes that she can utilize this force for the benefit of her brand.

However, she runs into a stumbling block when she sits down and starts to think about the

*Corresponding author
*Email addresses:* manuel.chica.serrano@gmail.com (Manuel Chica), wrand@umd.edu (William Rand)

plan logistics. Being part of a data-driven organization, she needs to answer a number of different questions to justify this marketing policy to her superiors. How does she balance this WOM program with her traditional marketing mix (Libai et al. 2013)? How much should rewards cost to maximize revenue (Ryu and Feick 2007) and how should she balance these costs with the number of customers who receive the reward (Schlereth et al. 2013, Stonedahl et al. 2010)? Which exact customers should she target (Hinz et al. 2011, Van der Lans et al. 2010)? Should she favor targeting influential users on social media (Watts and Dodds 2007, Trusov et al. 2010, Hinz et al. 2011)? Or should she only reward the highest revenue customers (Haenlein and Libai 2013)? Moreover, what will the effect of this promotion be on non-customers (Schmitt et al. 2011)?

This scenario is not fictional but one faced by managers on a regular basis. In fact, Blizzard Entertainment faced a similar set of questions when they sent statues of orcs from their massively multi-player online game, World of Warcraft, to all customers who had been playing the game for more than ten years (IGN 2015). Airline and hotel companies often consider these questions when designing special rewards for their loyalty programs (Terblanche 2015, Xie et al. 2015). And these questions are also relevant for new business models, such as freemium apps, where the area of interest centers around encouraging non-paying users to adopt paid contents (Kumar 2014). Of course, managers could spend time researching the previous theory and research on the various effects of WOM on product adoption and market expansion (Trusov et al. 2010, Libai et al. 2013), but what they really need is a practical tool that embodies this theory and provides them with direct answers to help them make decisions.

The goal of this paper is to provide a general framework for the creation of such a tool that can answer these questions. And we also provide an example of applying this framework to a real product, called Animal Jam, a multiplayer online freemium game for kids, created by WildWorks and the National Geographic Society. These questions about WOM are difficult to answer using traditional methods of analysis, due to the fact that WOM is fundamentally about interactions between customers, and is highly contingent upon the social network of the customers (Goldenberg et al. 2001). As a result, answering them requires an individual-level model that allows the analyst to represent the behavior of every customer, something difficult to do using many traditional forms of modeling.

A powerful solution to this problem is to use agent-based modeling (ABM) (Macal and North

2005, Epstein 2006), a computational approach where every individual can be represented separately and the entire context of their decision, including their social network and their adoption preferences, can be taken into consideration. The advantage of ABM is that researchers create the model at the individual level which does not require knowledge of higher level assumptions. Rand and Rust (2011) described a set of indicators to consider when deciding if ABMs are more appropriate than other tools such as analytical or statistical modeling. These indicators show that WOM-related marketing problems are effectively examined through the lens of ABMs. ABM is more appropriate than other quantitative tools when a complex and dynamic environment, such as a social network, is involved, and when the marketing measure of interest is an emergent result of consumer interactions, such as new conversions or revenue (Rand and Rust 2011). ABM also works well when the marketing research questions emphasize the heterogeneity of customers, and when the decision processes of those customers can be affected by different individual characteristics, seasonal behavior, media consumption, and the number and type of friends, with which they discuss the brand or the product.

Rand and Rust (2011) noted that the patterns of growth in the market that result from the interaction of many consumers are more complex than any individual's adoption decision. Because of this complexity, marketing studies are increasingly using ABM when analyzing new product growth (Garcia 2005, Delre et al. 2010), marketing adoption policies (Libai et al. 2013, Trusov et al. 2013), and targeting strategies (Haenlein and Libai 2013). Traditionally, many ABM models in the marketing literature have been used to advance marketing theory (Watts and Dodds 2007, Goldenberg et al. 2001; 2010), but ABMs can be calibrated using real data, and then can be used to develop insight into real-world applications (Stonedahl and Rand 2014). In this sense, a decision support system (DSS) can be built using ABM to help managers make real tactical and strategic decisions about marketing programs. The realism of ABM facilitates the understanding of the model and can make the DSS more comprehensible to stakeholders, since the model creates an ontology that is very close to the real world.

In this paper, we will explore a framework for creating an agent-based DSS to provide marketers and researchers with a new and powerful tool to help make WOM decisions and to better understand WOM phenomena. Using this framework we can situate customers (users) within a social network and give them their own individual states and actions (Wilensky and Rand 2015), and we can use this agent-based DSS to answer many of the questions of our brand manager from the beginning of this paper. Specifically, an agent-based DSS can assist marketers and managers to: (1) understand adoption dynamics and customer engagement (the way other people are affected by the engaged costumers, directly or indirectly, (2) leverage customer-to-customer interactions to improve business performance and (3), test and evaluate the effects of WOM and social value on revenue and product adoption in hypothetical market scenarios. The agent-based DSS can create market level outcomes by allowing the incorporation of individual behavioral rules (Libai et al. 2010) through a computational social network, which is representative of the real social network of costumers (Newman et al. 2006). These rules describe the typical activity of a customer and how they decide to adopt products or services by using diffusion information models (Rogers 2003) such as those based on cascade models (Goldenberg et al. 2001) and personal thresholds (Granovetter 1978, Watts and Dodds 2007).

Since the time of Little (1970), there has been a robust set of marketing models and DSSs for marketing (Wierenga et al. 2008, Lilien et al. 2013). Some DSSs were specifically built for modeling frameworks and decision-making processes in WOM (Lovett et al. 2013, De Bruyn and Lilien 2008), electronic online WOM activities (Dellarocas 2006, Cheung and Thadani 2012), and viral marketing (Van der Lans et al. 2010). This previous research shows that DSSs provide managerial benefits such as improving the decision making of marketers, enabling the exploration of more decisions, and updating the mental models of decision makers (Lilien 2011). However, many successful academic marketing models have a low level of practical use (Lilien 2011). Although models can produce significant benefits, many managers are reluctant to use them based solely on their objective quality in academic publications. Even in 1970, Little already noted that "most failures come from trying to deploy sophisticated, black-box optimization models in DSS environments because managers were unwilling to implement recommendations they did not understand". According to Lilien (2011), researchers must reduce the gap between the users' mental models and implemented decision models, which means helping DSS users understand and internalize the factors driving the model results and its managerial recommendations.

Therefore, our framework emphasizes the creation of an agent-based DSS that encourages participatory modeling to involve all stakeholders in the model creation process, and iterative modeling to incrementally build the model and discuss the results with all stakeholders (Voinov and Bousquet 2010). Moreover, our framework favors computational methods that facilitate the understanding of the models, and data-driven validation which allows both modelers and marketers to gain confidence in the DSS recommendations (Sargent 2005, Oliva 2003). In specific, our study presents a set of methodological guidelines, steps, and decisions to generate the models in the light of large datasets (Leeflang et al. 2015), with a clear focus on the managerial adoption of their results. A decisive step for ensuring this managerial adoption is its validation and testing as decision makers are often concerned with whether each model and its results are *correct* (Sargent 2005, Oliva 2003, Stonedahl and Rand 2014). Given the growth and availability of new data forms, we encourage modelers to follow a data-driven automated calibration process as the main validation tool and to use metaheuristics for automated model calibration (Miller 1998, Chica et al. 2016). We present the reasons why metaheuristics (Talbi 2009) are recommended here and the criteria and steps to design the most appropriate metaheuristic calibration method for each specific setting.

Additionally, we demonstrate the application of the agent-based DSS framework to a real hedonic freemium app, Animal Jam (*www.animaljam.com*). We follow the general guidelines and recommendations to construct the DSS and we show how to generate agent-based models using the app data. The Animal Jam DSS replicates and predicts premium adoptions in the app at both the macro- and micro-level. The DSS forecasts the number of new daily premium adoptions (macro-level simulation) and the specific users who are going to convert in a given time horizon (micro-level forecast). These two levels are useful for managers as they allow them to know the adoption rate and evolution of the number of premiums in a medium-term but also to compute the likelihood of a basic user to convert in a short-term horizon taking into account her/his friends.

Within this application, we examine diffusion mechanisms such as an agent-based version of the Bass model (Bass-ABM) (Rand and Rust 2011, Bass 1969) and the complex contagion model of Centola and Macy (2007). To our knowledge, this is the first work in marketing where adoption dynamics are modeled through using the complex contagion model. We will discuss the complex contagion model in more detail later, but the basic idea is that adoption is contingent primarily on the absolute number of people you know that have adopted. We also use a metaheuristic automated calibration tool to tune the parameters of the models with respect to the historical data provided by the firm. Later in the paper we use the validated model to explore targeting and referral marketing policies (Schmitt et al. 2011, Haenlein and Libai 2013) and evaluate their impact to expand the premium market through WOM and customer engagement for the specific Animal Jam setting.

## FRAMEWORK AND STEPS TO BUILD A DECISION SUPPORT SYSTEM

Our framework presents guidelines, design steps, and specific recommendations for creating an agent-based DSS for WOM market scenarios. The basic foundation for the DSS framework that we propose is an individual-level model which captures the social interaction dynamics of the customers, as embedded in a social network. Figure 1 shows the main four guidelines and three main steps to consider when building the agent-based DSS. The four guidelines are: ($G1$) to follow an iterative and participatory modeling process with marketers and stakeholders, ($G2$) to analyze and use all the available data to build the DSS, ($G3$) to employ data-driven calibration to increase users' confidence, and ($G4$) to minimize the complexity and the number of parameters in the model in order to increase ease of understanding. These guidelines are not necessarily steps to be followed in order, but rather guiding principles to keep in mind during the construction of the DSS that will maximize the probability of its adoption by marketers. We will describe the specific steps for creating the agent-based DSS in a few paragraphs.

The first guideline $G1$ encourages the modeler to follow an implementable iterative model-building process (Leeflang et al. 2015) with a participatory element. This process must involve marketing managers and stakeholders, which is a key ingredient in facilitating better decisions, with less conflict and more success (Voinov and Bousquet 2010). $G1$ is also related to the *trialability* aspect of innovation adoption (Rogers 2003), since if marketers and stakeholders are able to try out the DSS, they are more likely to adopt its use consistently in practice.

Given the growing importance of digital data for companies, guideline $G2$ discusses how all available data should be used when constructing the DSS. This data serves as an input for the creation of the DSS and may include information about the real social network of potential customers, seasonal information about customer use, empirical data on product or service adoption, and WOM volume or sentiment.

Guideline $G3$ states that data-driven model calibration is the cornerstone of the validation of the agent-based DSS. This validation is vital to increase the confidence of stakeholders in the DSS recommendations (Sargent 2005). The modeling process of the DSS is aimed at enhancing

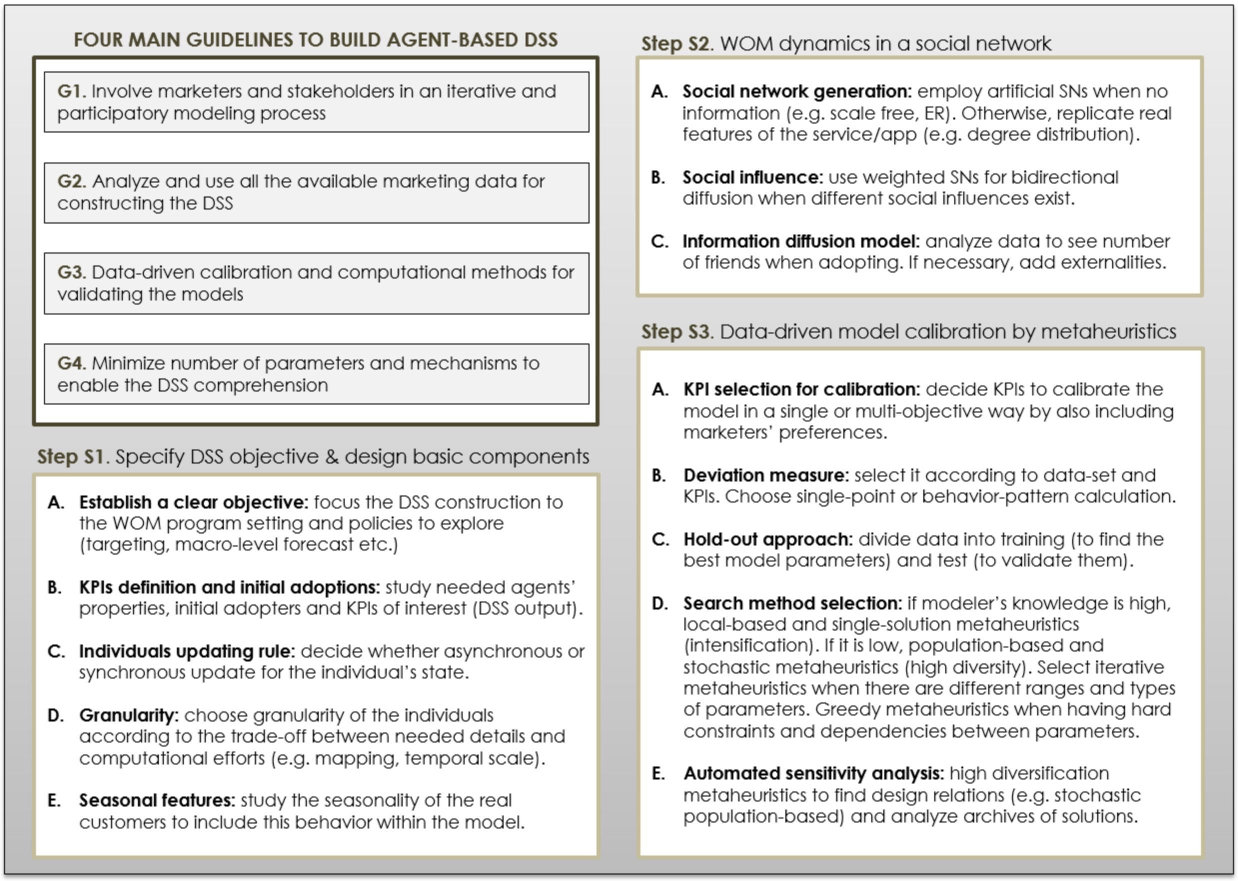


Figure 1: Four guidelines and three steps to build an agent-based DSS for WOM programs.

marketers knowledge and understanding of the WOM dynamics by identifying the impact of the solutions and supporting marketing decisions for the WOM program. It is also important for the model to be understandable to encourage stakeholder use of the DSS. Therefore, guideline $G4$ declares that modelers should use the minimum number of parameters and minimum number of mechanisms that enable satisfactory and valid results (Terano 2008) within the DSS. By creating minimal models the researcher is more likely to facilitate the understanding of the models by stakeholders, and at the same time is more likely to create a valid model. As Axelrod (1997) explained, *keep it simple, stupid*.

Taking into account these four guidelines, we describe a framework for the creation of the DSS, based on three main building steps, which are summarized in Figure 1. These steps are to: ($S1$) specify the marketing objective and the basic components of the DSS, ($S2$) create the model of the WOM dynamics through a social network, and ($S3$) use metaheuristics to perform a data-driven calibration of the model. In the following subsections, we will explain each of the steps in turn.

*S1. Specify the Decision Support System Objective and Design the Basic Components of the Agent-Based Model*

To understand WOM dynamics we need to specify a framework upon which to build the agent-based DSS and to embed the adoption dynamics. We recommend the use of ABM (Macal and North 2005, Epstein 2006) because it can effectively model the aggregate consequences of WOM on the basis of local interactions among individual members of a population (Libai et al. 2013, Goldenberg et al. 2001). Web Appendix A goes into considerable detail about why ABM is appropriate as the basis of the DSS, but in brief, ABM provides the ability to model a large number of heterogeneous individuals interacting across a complex social network where the agents take their own actions that affect their decisions about how to spread WOM. During this construction step $S1$, five sub-steps should be taken to define: ($S1A$) the objectives of the system, ($S1B$) the model architecture, ($S1C$) the updating behavior of the agents, ($S1D$) the granularity of the agents, and ($S1E$) seasonality of user behavior. Though designing a model architecture is important for any DSS, we focus here on the decisions points that should be considered by someone interested in developing a DSS for a WOM program.

**S1A: Establish a clear objective.** The first step when building the agent-based DSS is to keep in mind the main marketing objectives and the potential WOM programs that the stakeholders would like to explore. This is based on the design principle of building the model toward the question that the model is meant to answer in an incremental fashion (Wilensky and Rand 2015). Additionally, the intended use of the models should be defined as precisely as possible (Leeflang et al. 2015) and every decision should be made with this in mind throughout all of the building steps of the DSS. In many cases, managers start with a default question along the lines of which users should I incentivize in order to maximize the adoption of my product?

**S1B: KPIs definition and initial adoptions.** Agent-based DSS are discrete-time simulations which end after several time steps. When a time step ends, the simulation collects the key performance indicators (KPIs) of interest at each time step and returns them as the output of the DSS (e.g., the time-series of product-service adopters). Typically, in many WOM campaigns, the basic assumption is that some users have already adopted the product or the service (De Bruyn and Lilien 2008). This is usually modeled through the use of a binary state, i.e., an agent is labeled as a *non-adopter* or *adopter*. Then, it is necessary to define how these initial adoptions are chosen. It could be through random choice, or potentially influenced by some empirical data. A default option is to start by defining a single KPI (e.g., number of purchases or total adoptions) and then running the model to examine the adoption of the product in a baseline condition without any WOM program. Afterwards, one can see how incentivizing users affects the adoption rate.

**S1C: Individual updating rule.** The individuals of the agent-based DSS can act asynchronously or synchronously within the simulation. Synchronous updates occur when no individual reveals their new state until all individuals have had a chance to change their state. This is usually set to occur during a particular system-level event which represents the time step of the model. Asynchronous updates occur when individuals act and immediately reveal their state (Wilensky and Rand 2015). Synchronous updates are useful if there is not constant interaction between individuals involved in the adoption process, but rather some time lag between when they adopt and when the information about adoption can be passed on to others. This could be the case if there are no signs of conspicuous consumption for instance, and the adoption decision only becomes obvious once users' discuss their adoptions. Asynchronous update makes the simulation closer to traditional continuous diffusion models such as the original formulation of the Bass model (Bass 1969) because an asynchronous update is closer to assuming that the number of adopters at time $t$ not only depends on the number of adopters at time $t-1$ but also on the instantaneous number of adopters at time $t$ (Rand and Rust 2011). As a result, the default rule is to use asynchronous updates for most agent-based DSS.

**S1D: Granularity.** An important question is to decide the granularity and mapping of the individuals within the model. This requires specifying the temporal scale of the market context and the number of real customers represented by an agent within the model. At one extreme this could be modeling every real customer with exactly one agent, and at the other extreme it could mean representing thousands of customers with one agent. This decision is made based on the needed granularity of the DSS to make decisions. For instance, is it necessary to make decisions about individual behavior, or is segment or population-level behavior a good enough representation? In addition, the computational resources needed to run the model should be taken into consideration (Wilensky and Rand 2015). The default rule is to find a good trade-off between these two factors by always representing the fundamental level of information, necessary to answer the WOM marketer's questions. For instance, if the question is about which exact individuals to incentivize then a one-to-one mapping is often needed.

**S1E: Seasonal features.** The heterogeneity and flexibility of agent-based DSSs permit an easy inclusion of seasonal patterns of behavior. As seasonality affects product adoptions (Peers et al. 2012, Guidolin and Guseo 2014), it is important to model seasonal effects when constructing an agent-based DSS. The modeler can define seasonality effects in production acquisition, service usage, or digital access at a given time step. For instance, the modeler can define a time-varying parameter that controls the probability of a particular event occurring at any particular time. At each time step in the simulation, the DSS can first consider if a customer takes an action, such as accessing the service or using the product, by drawing a random number from a uniform distribution and by comparing it to these seasonality parameters. The default suggestion for this step is to analyze the data and base the seasonality processes directly on this data.

*S2. WOM Dynamics in a Social Network*

Social networks play a fundamental role in the way information reaches consumers, channel members, and suppliers (Goldenberg et al. 2009, Van den Bulte and Joshi 2007). This is because the individual adoption decisions of customers normally depend on two factors: (a) external influence (by salespeople, advertising, promotions, and other marketing efforts) and (b) internal influence (affected by WOM or by observing conspicuous consumption of someone in their social network) (Libai et al. 2010). A social network is generally defined by a set of actors and the relationships (ties) among them. The social network properties of an individual can impact the success of marketing actions, such as pricing or promotion strategies (Godes and Mayzlin 2009). Within this step $S2$ of our framework we define three sub-steps: ($S2A$) generate the social network structure, ($S2B$) define the social influence between individuals, and ($S2C$) model how information dynamics occur in the social network.

**S2A: Social network generation.** The social network defines the relationship between different consumers or users. Although it is a common approach to approximate a real social network with a synthetically generated preferential attachment network (Barabási and Albert 1999), many studies have provided evidence that most of the real-world social networks have distinct structural properties from these synthetically generated social networks (Newman 2003, Stonedahl et al. 2010). Given the growth of online data and social media platforms, at least partial information and data about the social networks used by customers frequently exists about the customer base. Whenever possible, the DSS should be designed to take into account all the existing information about the real social network of the marketing context ($G2$ guideline). Therefore, the default rule of this step is that, if the actual social network is known, then that should be used. Otherwise, the modeler is encouraged to use all the available information to guide the network generation by replicating the properties of the real-world social network (e.g., degree distribution, average degree, or cluster coefficient). If no social network properties are known, the marketer can use one of the synthetic networks that exist, but (s)he should consider whether it is worth creating a DSS for WOM if no social network information is available, as the social network is a critical component of understanding WOM programs.

**S2B: Social influence.** Many standard social network models assume that all consumers exert the same influence on each other, but in reality we know that is not the case. It is useful to consider the role of social influence between the WOM actors. One typical approach is to enrich the WOM process by modeling heterogeneous social influence by a weighted social network. To do this, the links of the social network are weighted based on the social influence between customers. If all the links have the same weights, then the model collapses to the traditional model with homogeneous influence. The default rule is to always include these social influence differences which can be inferred from data analysis after the implementation of marketing social-based activities by the company. For instance, it may be possible to observe when two users are talking with each other and use that information as a model of social influence. When not available, modelers can first equally set these weights for all the customers and later, run a sensitivity analysis to evaluate the effect of changing the weights.

**S2C: Information diffusion model.** Apart from designing the network structure, it is important to model the WOM dynamics that occur through the medium of the network. The probability that a particular individual chooses between one product alternative or another is increased according to the relative number of others choosing the same alternative and the influence of those others on the focal individual (Watts and Dodds 2007, Goldenberg et al. 2001, Trusov et al. 2013, Van den Bulte and Joshi 2007). Threshold (Granovetter 1978, Watts and Dodds 2007) and cascade models (Goldenberg et al. 2001; 2010) are common individual-level diffusion mechanisms used in marketing for modeling an individual's decision. Stochastic cascade models hypothesize that when each individual adopts a product they have a small probability of influencing any of their social neighbors to adopt the product. In the threshold model, each individual observes the fraction of neighbors that have adopted and then adopts if this fraction exceeds a certain threshold. In both cases, it is also possible to add an external influence parameter, which can encourage adoption of the product independent of social influence. Unless data analysis suggests a more appropriate approach, it makes sense to start with something similar to an agent-based version of the Bass model (Bass-ABM) (Rand and Rust 2011), which assumes independence of the internal (customer-to-customer interactions through WOM) and external effects (Libai et al. 2010). This model is a form of a cascade model, and is already well-accepted in the marketing literature.

*S3. Data-driven Model Calibration by Metaheuristics*

Automated calibration is a data-rich and computationally intensive process that uses an error measure to compare real-world data to model-data, and then tunes the parameters of the model in order to identify a set of parameters which best match the data (Oliva 2003, Sargent 2005, Stonedahl and Rand 2014). Automated calibration attempts to discover the best parameters of the model that fit the model to the data. This evaluation of the model fitting is done by running the computational model and comparing its outputs to the data. This means that automated calibration requires an error measure and an optimization method to modify the parameters in a systematic way in order to minimize the error measure. Our framework presents the building steps to calibrate the agent-based DSS that was created in the previous two steps, using metaheuristic methods (Talbi 2009).

Metaheuristics are a family of approximate non-linear optimization techniques that provide *acceptable* solutions in a reasonable time even when problems are hard and complex (Talbi 2009). When calibrating a complex system such as an agent-based DSS, metaheuristics are preferred compared to gradient-based methods or mathematical programming for two main reasons: (1) we can only make minimal assumptions about the non-analytical simulation model (i.e., the relationship between all parameters and all outputs in the ABM framework is unknown), and (2) the objective function (i.e., a function that formulates the goal to achieve when optimizing) is time-consuming and needs to be run many times in order to accurately compare the simulated model to real marketing data. This makes it difficult to create a closed-form solution, and computationally too expensive to conduct a full search of the parameter space.

One well-known type of metaheuristic is the genetic algorithm (GA) (Goldberg and Holland 1988), which are powerful search methods that have already been applied to marketing problems (Luo 2011, Venkatesan et al. 2004). Metaheuristics can be classified according to various characteristics (Talbi 2009): nature-inspired versus not nature-inspired, deterministic versus stochastic, population-based versus single-solution based search, and iterative versus greedy. But the relevant issue when building an agent-based DSS is to find the most suitable metaheuristic for the particular WOM setting, which can be difficult, since there are two contradictory goals that must be taken into account when choosing and designing the metaheuristic. On the one hand, there is a goal of maximizing the exploration of the parameter space (diversification) and, on the other hand, there is a goal of exploiting the best solutions discovered so far (intensification). In the next five steps we will discuss these and other criteria to select the most appropriate metaheuristic calibration method.

**S3A: KPI selection for calibration.** The first step is to identify one or more KPIs to compare the output of the DSS with real data. These KPIs can be the number of adoptions, WOM volume and/or sentiment, or the sales of a brand. Depending on the number of KPIs in conflict, modelers must choose between single-objective (only one KPI) and multi-objective metaheuristics (when more than one KPI needs to be optimized at the same time and they are potentially in conflict) (Talbi 2009, Chica et al. 2010). An additional option when more than one KPI exists is to include stakeholders knowledge to weight and value the KPIs within the calibration process by using preference relations or units of importance between the defined KPIs (Chica et al. 2011). This enables the creation of a single KPI by taking a weighted sum of multiple KPIs. However, in the spirit of keeping things simple, the best place to start is calibrate the DSS with a single KPI and later, move to a more advanced approach, if needed.

**S3B: Deviation measure.** The metaheuristic deviation measure evaluates the quality of a set of parameters by comparing the model results for that set of parameters to historical data. A modeler can use different error or deviation measures and this choice can significantly affect the calibration performance (Stonedahl and Rand 2014). A traditional approach is to use single point-based measures (e.g., root mean square error (RMSE), Euclidean distance, or mean absolute percentage error (MAPE), among others). The selection of the specific measure depends on the problem and data characteristics (e.g., trends in historical data or number of KPI datasets). The default rule of this step is to use RMSE or Euclidean distance to calibrate a simple set of KPIs for a match of a typical series of historical data points. If the goal is to favor general trends over specific data matches, then we suggest using MAPE since it decreases the effect of big isolated data point errors within the model calibration (Chai and Draxler 2014).

**S3C: Hold-out approach.** Independent of the choice of metaheuristic, modelers are encouraged to use a hold-out approach when calibrating the model. As explained in Stonedahl and Rand (2014), the modeler divides the whole historical dataset of KPIs into two datasets: training and test, with their corresponding environmental variables. The environmental variables are those variables that will not change from run to run of the models as they do not belong to the set of parameters to be calibrated. The calibration of the model is then accomplished by identifying, using the training dataset, the model parameters that minimize the chosen deviation measure. Later, the model's validity is examined by using the same set of parameters but for the test dataset. More advanced approaches such as k-fold cross-validation can be also used (Witten and Frank 2005). By default, the rule is to always use, at least, a basic hold-out approach to make sure the calibrated model is generalizable to data that was not used to train the model.

**S3D: Search method selection.** The knowledge of the modeler about possible good parameters' values and the features of the parameter space affect which metaheuristic to use. The modeler knowledge should be used to help constrain the search space and help create a metaheuristic that can really explore the restricted parameter space in-depth, a process called intensification (Talbi 2009). When knowledge about the best configuration parameters is scarce, population-based metaheuristics enable a wider search for the best set of parameters' values. Additionally, iterative metaheuristics (i.e., those such as GAs that start with a complete solution or population and transform it at each iteration using search operators) are more flexible if different parameter ranges and types exist when calibrating (e.g., integer, real, or binary). This is because the solution is built at the start and there is not a constructive process that needs to be customized for each type of parameter (Talbi 2009). Greedy and constructive metaheuristics (e.g., GRASP (Feo and Resende 1995) or iterated local search (Lourenço et al. 2003)) are suggested when there is a high number of dependencies between parameters and hard optimization problem constraints (Chica et al. 2010). The main reason is a greedy metaheuristic starts from an empty solution and constructively assigns at each step a parameter value for the calibration problem until a complete solution is obtained.

As a default recommendation, when modeler knowledge is limited it makes sense to use as a default one of the population-based metaheuristics, such as the previously mentioned GAs (Holland 1975), ant colony optimization (Dorigo et al. 1996), or particle swarm optimization (Kennedy 2010). In general, we recommend iterative metaheuristics given the unusual presence of hard parameter constraints in this kind of DSS.

**S3E: Automated sensitivity analysis.** Sensitivity analysis reveals those parameters to which the model behavior is highly sensitive (Saltelli et al. 2008). Together with calibration, it is a key ingredient for model testing and verification (Oliva 2003). Miller (1998) initially pointed out GAs as an appropriate tool for sensitivity analysis because of their capability to explore a wider range of parameter settings with a higher resolution and to also consider potentially complex or non-linear interactions between them. Specifically, population-based metaheuristics offer not only a final calibration solution but an *archive* of evaluated model calibration solutions for the model. Automated sensitivity analysis can be performed on these solutions to discover hidden properties related to the model design (Chica et al. 2016). The default rule is to use metaheuristics that emphasize diversity, often in the form of stochastic population-based metaheuristics, to assist in the sensitivity analysis of the model.

## APPLICATION TO A FREEMIUM BUSINESS MODEL

Freemium business models offer a service or a product free of charge but a premium is charged for advanced features, functionality, or related products and services (Anderson 2009). Over the past decade, freemium has become an important business model for digital-based products and services including software, games, and websites (Teece 2010). The 2015 freemium app monetization report of App Annie (an app analytics firm) & IDC (International Data Corporation) stated that freemium app revenues grew by over 70% while paid app revenues declined by 19% from 2013. WOM often plays a large role in freemium app adoption, since many of these apps are game-related or have a built-in social network component Bapna and Umyarov (2015). We will examine in this section the application of our agent-based DSS framework for WOM to one real-world, hedonic, freemium app.

The importance of WOM in freemium models is not new. Cheng and Tang (2010) showed that the optimal price of commercial software increases with the network intensity of the software, which can be enhanced by free trials. Oestreicher-Singer and Zalmanson (2013) studied the premium services of Last.fm, an online music website, and found that the willingness to pay for premium services is strongly associated with the level of community participation of the user, i.e., social behavior. Bapna and Umyarov (2015) showed that peer influence caused more than a 50% increase in the probability of adopting premium services due to the influence coming from an adopting friend. They also found that users with a smaller number of friends experienced stronger relative increase in the adoption likelihood due to influence from their peers as compared to the users with a larger number of friends. Other studies presented relevant conclusions on the importance of WOM programs for managing freemium business models. For instance, Lee et al. (2013) examined consumer referral behavior and showed that referral invites recruit consumers who later convert to premium consumers. Therefore our paper makes a vital contribution to this literature by illustrating how a general DSS for WOM can help managers gain insight into many of these questions for their particular datasets that have previously been explored by marketers.

*App Data and the Main Marketing Questions*

In this section, we will explore the specific application of our framework to a freemium app called Animal Jam (*http://www.animaljam.com*). It is a multi-platform social online game for kids where users interact online with other users. Basic users can freely access and play with the game and interact with other users both premium and basic, but premium users receive additional benefits such as weekly in-game currency allowances (called *diamonds* and *gems*), the ability to adopt virtual pets, access to all of the avatars, and premium-only adventures. WildWorks, the developer of Animal Jam, was interested in making better marketing decisions for expanding their premium market. They wanted to understand the real importance and role of WOM when free users adopt premium services (conversions). They were also interested in examining the efficacy of marketing campaigns among their current customers to reward and target them with gifts and bonus packages to incentivize those customers to talk positively about premium membership.

Generally, *referral rewards* are designed to motivate consumers to spread positive WOM about products and services and to turn customers into an element of the sales force (Biyalogorsky et al. 2001). Lee et al. (2013) examined rewarding through referrals in the context of a freemium software service. They explored the question of the right referral bonus incentive to offer to free users. Contrary to the belief that more is better, they found, by maximizing the average consumer referral rate and changing the referral incentives, that there exists an optimal incentive point for referrals that is not simply as much as possible. Clearly then the decisions related to how to implement an optimal referral program are not straightforward. In our case and to help Animal Jam managers with their marketing goals, a DSS was constructed using the guidelines and steps described above, in order to: (1) replicate and forecast the conversion rate from freemium to premium members, and (2) evaluate incentivization-based marketing campaigns and measure the additional customer acquisition created by *amplified WOM*.

The company provided us with daily conversion and subscription data of their users in 2012. This data included 1.4 million game users and their social network to analyze and is valuable as it gives us the opportunity to properly design the DSS and to validate the models with respect to historical trends. For the quantitative models' validation we restricted the whole data of 1.4 million users to only those ones active during the months of the study (i.e., from 40,000 to 50,000 basic and premium users).

In the following sections, we will explain how we built a DSS for Animal Jam utilizing our framework to help answer the main questions of interest to managers. Note that, in addition to the detail provided in this section, we present the documentation, verification, and validation details of the DSS for Animal Jam in Web Appendix B. Furthermore, the reader can freely access the agent-based DSS source code at *https://bitbucket.org/mchserrano/socialdynamicsfreemiumapps*.

*S1: Definition of the DSS Objective and Designing Basic Components*

During this step, we paid close attention to guideline $G2$ by using the data provided by the company and to guideline $G4$ by simplifying the DSS design.

**S1A: Establish a clear objective.** The main goal in building the Animal Jam DSS is to forecast the adoption of the premium services by users of the app. Since the incentive programs of managers' interest often focus on individual-level policies (e.g., targeting specific users) we should consider a DSS based on a very granular model. The DSS should be able to forecast both the total number of new daily premium adoptions of the app (macro-level forecast) and whether a specific user is likely to become a premium user in a given time horizon (micro-level forecast).

**S1B: KPIs definition and initial adoptions.** There is one central KPI: the number of premium adoptions of the basic users of the app. This KPI has to be forecasted daily or weekly as a conversion rate for the total number of users of the app. As noted by guideline $G1$, an iterative process with stakeholder participation was followed here to find the correct KPI, necessary to answer the WOM marketer's questions. During this process, we discovered there was an interest in not only forecasting the aggregate conversion rate, but also in forecasting particular individuals who were likely to convert in the near future. To begin to calibrate our model we defined an initial adoption rate ($\alpha$) which reflects the number of initial premium adopters of the app. The simulation of the DSS will then start with a ratio $\alpha = 0.0406$ which corresponds to the 4.06% of premium users of the real app at the starting simulation date.

**S1C: Individual updating rule.** We used an asynchronous update rule in the model to better replicate the WOM dynamics of the app, following the default rule of this framework step. Within the same time step (day), the individuals can make decisions and update their states asynchronously. This more closely mirrors the real app, where as soon as someone becomes a premium member they can influence their neighbors who are online.

**S1D: Granularity.** After a preliminary exploration of granularity, we defined one individual of the ABM framework to correspond to two real app users active during the months of the forecast. Therefore, our simulation consists of 20,000 individuals (with 812 of them being initial premium adopters) that map to 40,000 active real users of the app during the months of the study. This assumption was made in order to decrease the computational cost of the model without giving up too much resolution.

**S1E: Seasonal features.** We found a strong seasonality in the app dynamics after analyzing the time series of premium conversions in the historical data (guideline $G2$). The app seasonality shows that during weekends people tend to have more activity than on weekdays. The observation seems obvious as it is a hedonic app, aimed at kids, normally played by users when they have more free time. Please note that the seasonality of the Animal Jam DSS is crucial as app users cannot adopt premium content and/or talk with their direct friends on the app if they do not access the app on a particular day. We define seasonality within the model by grouping the probability of using the app in to two seasonality parameters: the probability of using the app on a weekday ($\gamma_0$) and the probability of using the app on a weekend ($\gamma_1$). This simplification is in line with guideline $G4$ to keep the model as simple as possible without losing accuracy (Terano 2008, Wilensky and Rand 2015).

*S2: Replicating the WOM Diffusion Process through a Social Network*

The analysis of the freemium app data shows there is a clear difference in friendship patterns between basic and premium users, observed in the degree density graphs (guideline $G2$). The fact that premium members show such a different pattern of social behavior lends credence to the related literature on freemium business models that show that the social activity of users is related to their conversion (Bapna and Umyarov 2015, Lee et al. 2013). Further supporting the importance of social effects is the observation that having more premium members as friends increases the likelihood that one will adopt. Specifically from analyzing the app data we see that (a) less than 1% of users without a premium friend subscribe, (b) the chance of conversion triples (from 0.8% to 2.3%) if a user goes from having zero premium friends to just one, and (c) the incremental impact on adoption of having premium friends is higher the fewer friends you have. Using these insights and the $S2$ step of the framework we constructed a WOM diffusion model for the DSS, which we will explore in the next few paragraphs.

**S2A: Social network generation.** By analyzing the real social network we obtained a distribution degree which represents the social relations of the app users. One interesting fact about this distribution degree is that the app limits the number of friends any user can have to 100 friends and, as a result, the degree distribution is heavily bi-modal. As a result, there is a group of users who have very few friends and another group clustered around the upper limit of 100 friends, which is enforced by the Animal Jam app, with fewer users in between these extremes. Using this information we employed the *generalized random networks* algorithm of Viger and Latapy (2005) to generate a social network of individuals with similar distribution degree to the real social network of users in the real app. Though we have the actual network information of the model, and could have used that, by using a synthetic network, we are able to scale the size of the network to the size of the population being simulated. Web Appendix B provides more information about the degree distribution of the generated social network.

**S2B: Social influence.** Following the $S2$, for each link between two individuals $i$ and $j$ of the social network, two different weights, $\tau_{i,j}$ and $\tau_{j,i}$. These two weights generate a bidirectional influence, where the social influence of $i$ when (s)he talks with $j$ can be different than when $j$ talks to $i$. When $\tau_{i,j} > k$ (or $\tau_{i,j} < k$) the influence of $i$ on $j$ is increased (or decreased) according to the weight. Initially for the model and because we do not have specific information about the social influence between the users, we set the influence $\tau_{ij}$ to 1 for all the links. However, we did investigate using communication records from the real game to model social influence, but in this case, it did not appear to improve the results.

**S2C: Information diffusion model.** Our data analysis and existing work on freemium apps (Bapna and Umyarov 2015) highlights the marginal impact of additional premium friends when a user already has a large number of premium friends. This fact suggests the use of a complex contagion model (Centola and Macy 2007) as the diffusion method of the agent-based DSS. The model is called a *complex contagion* model because successful transmission depends upon interaction with multiple carriers. Complex contagion is close to the threshold model but the main difference is that in the threshold model, the threshold $\phi$ describes the exposure amount necessary to convert relative to the number of friends an individual has, where in the complex contagion model the threshold $\phi$ is the absolute (not relative) number of exposures necessary to convert. We also implemented an agent-based Bass Model (Bass-ABM) (Rand and Rust 2011)). The Bass-ABM translates the hazard rates of the Bass model to probabilities for a single consumer, and embeds the consumer in a social network, where the decision to adopt also depends on the fraction of neighbors that have converted. Finally, an extended complex contagion model is considered. This new diffusion model adds an innovation probability to treat diffusion similar to the Bass-ABM, which explicitly includes an innovation probability. Web Appendix B provides more details about these three diffusion models: Bass-ABM, traditional complex contagion model, and extended complex contagion model having an innovation probability.

*S3: Designing the Data-Driven Calibration Method*

We describe here how to design the automated metaheuristic calibration for the DSS of the Animal Jam (step $S3$ of the framework). This step is used to validate the model (guideline $G3$) (Sargent 2005). The parameters of the agent-based DSS must be calibrated to adequately fit the historical data of premium adoptions. The size of the parameters set $P^*$ to be calibrated is either three or four, depending on the diffusion model: two probability values for seasonality ($\gamma_0, \gamma_1 \in [0, 1]$), and the innovation and imitation coefficients ($\hat{p}, \hat{q} \in [0, 1]$), or the minimum threshold of premium friends for the complex contagion ($\phi \in [0, m]$, where $m$ is the maximum number of friends). Table 1 shows the list of parameters to be calibrated for all the considered diffusion models. The decisions necessary to design the most appropriate metaheuristic calibration method are discussed in the following paragraphs (step $S3$).

| **Diffusion models of the DSS** | **Number of parameters** | **Seasonality parameters** | | **Bass-ABM coefficients** | | **Complex contagion thresholds** |
|---|---|---|---|---|---|---|
| | | $\gamma_0 \in [0,1]$ | $\gamma_1 \in [0,1]$ | $\hat{p} \in [0,1]$ | $\hat{q} \in [0,1]$ | $\phi \in [0,m]$‡ |
| Bass-ABM | 4 | ✓ | ✓ | ✓ | ✓ | - |
| Complex contagion | 3 | ✓ | ✓ | - | - | ✓ |
| Complex with innovation | 4 | ✓ | ✓ | - | ✓ | ✓ |

‡where $m$ is the maximum number of friends an individual can have in the app.

Table 1: Set of model parameters $P^*$ to be calibrated by the GA.

**S3A: KPI selection for calibration.** We use a single KPI of interest, the aggregated number of new premiums (adoptions) per day in the app. We could instead use individual conversions as the KPI but it turns out that calibrating using the aggregate patterns also creates an accurate tool for micro-level targeting as we will explore in the validation section of the experimentation. The quality of the calibration solutions generated by the metaheuristic are evaluated by calculating the difference between the historical new premium users and the model output in a daily setting. This model output is obtained by running a Monte-Carlo (MC) model simulation of 15 runs for each parameters setting, which was sufficient to account for the model's variability.

**S3B: Deviation measure.** A single-point approach is followed to calculate the deviation between the real and simulated new premium values. We use the Euclidean distance as the objective function of the metaheuristic after running a preliminary experiment to check the goodness of fit to the seasonal trend using this measure.

**S3C: Hold-out approach.** A hold-out approach over the whole dataset is carried out as recommended by the default rule of the framework. From a dataset of three months of daily conversions, two datasets are generated: 60 days of historical premium daily conversions as a training set and 31 days for the test dataset. Therefore, the metaheuristic calibrates the models' parameters using the training data to fit historical data and then, this calibrated model is used to predict premium conversions on the test dataset in order to validate the model generalization.

**S3D: Search method selection.** We did not have any particular model knowledge for the set of parameters we were exploring. Additionally, there is not any hard constraint on the relationship between the parameters and each of the three diffusion models has a set of parameters of different types. For instance, threshold $\phi$ of the complex contagion model is an integer parameter while the seasonality parameters $\gamma_0$ and $\gamma_1$ are real. Therefore and following the default rule of this framework step, we used a high diversity, iterative, and population-based metaheuristic (e.g., a GA (Holland 1975, Goldberg and Holland 1988)).

**S3E: Automated sensitivity analysis.** We are also interested in understanding the ranges of each parameter that seem to provide the highest validity and the actual forecasts for the freemium app. For instance, the values for the threshold parameter of the complex contagion model can provide interesting insights about the how many freemium friends a focal friend needs to convert. The use of a population-based metaheuristic has the advantage of performing an automatic sensitivity analysis since it generates a large set of solutions to the problem (Chica et al. 2016). This enables the exploration of the already evaluated set of parameters and allows the examination of non-linear interactions between the parameters.

The recommended metaheuristic characteristics given by the latter steps suggest that a GA can work well when calibrating the DSS for the Animal Jam app. In a nutshell, a GA evolves a population of solutions (chromosomes) each of which represents a model's parameter set. These solutions are evolved until achieving the best possible design (parameters of the model) for a given modeling goal (i.e., correctly matching historical data). We detail the remaining GA components and the GA parameters in Web Appendix C. Finally, it is important to remark that the designed GA calibration method is independently run to calibrate each of the the three diffusion models: Bass-ABM, original complex contagion, and complex contagion with innovation coefficient. We can then use the accuracy of the macro-level forecasting of the premium adoption of the app to adjudicate which model does the best job of describing the real WOM behavior.

## VALIDITY AND PERFORMANCE OF THE MODELS FOR THE FREEMIUM APP

Modelers and managers need to know how well the DSS models replicate the real-world data after the calibration. The agent-based DSS for the app case provides two different forecast levels: macro- and micro-level forecasts. We first analyze the agent-based DSS to examine its ability to replicate macro-level historical data, i.e., the number of daily new adoptions provided by the calibrated models. Later, we examine the models' performance when forecasting the most likely users to adopt premium in the near future (micro-level forecast).

*Model Calibration Results for Macro-level Forecast*

The GA run for calibrating the DSS for Animal Jam ends when $20,000$ different parameter settings (solutions) have been tested. This is the stopping criteria of the GA and is sufficient for finding a good quality solution. We also run the overall GA calibration method 15 times per model because the GA itself is non-deterministic. At the end of all the runs, the metaheuristic calibration method returns the mean and standard deviation values of the Euclidean distance between the real and simulated data.

The Bass-ABM presents the best fitting within the set of three diffusion models of the DSS. The GA calibration of the Bass-ABM, ends with a Euclidean mean value of 358.72109 and a standard deviation of 2.00986. By following the hold-out approach, we examine the calibrated Bass-ABM model on the test dataset by obtaining an Euclidean mean value of 339.01315 and 9.09134 for the standard deviation. The DSS with a complex contagion model obtains mean values of 447.15797 and 420.53968 for the training and test datasets, respectively. The standard deviation is again higher for the test dataset: 20.37237 for test while obtaining 4.75522 for training dataset. Finally, the complex contagion variant with the innovation coefficient obtains similar results to the original complex contagion model. Its Euclidean mean values are 441.52756 (training) and 416.76062 (test) while its standard deviation values are 5.45638 (training) and 19.4957 (test).

By analyzing these results we observe how the Bass-ABM fits the historical data better than the complex contagion models both in training and test datasets, although the differences are low. Differences between both complex contagion models are also low. The low standard deviation values indicate that the 15 GA runs were sufficient to obtain quality results. Therefore, these macro-level results of premium adoptions for the Animal Jam setting suggest that the introduction of an innovation coefficient $\hat{p}$ does not improve the performance results of the traditional complex contagion model. Web Appendix D provides an additional analysis of the temporal evolution of premium adoptions.

By using the population-based results we can directly perform a sensitivity analysis and observe the parameters' range to better understand the strength of the social influence when adopting premium services. Web Appendix D also presents boxplots of the parameters' distributions for the three diffusion models. From this sensitivity analysis we can see that the threshold parameter of both complex contagion models ($\phi$) is mostly set to three friends by the calibration process. In the case of the complex contagion model with innovation probability, the range of good values for the innovation coefficient ($\hat{p}$) is much wider than in the Bass-ABM. This fact means that the influence of the innovation coefficient for the complex contagion model is less critical than for the Bass-ABM and also reinforces the above-mentioned conclusion of low differences between both complex contagion models. Finally, we notice that the GA-based calibration method focuses on very specific seasonality and imitation coefficient values to provide a good model fit for the complex contagion model. On the other hand, the Bass-ABM results are more robust with respect to a wider range of values for the imitation coefficient and seasonality parameters ($\hat{q}$, $\gamma_0$, and $\gamma_1$).

*Micro-level Forecast: Models Comparison*

One of the main interests of the managers was a better understanding of how to target and incentivize basic users. The objective was to identify those users who could easily be convinced to become premium users. To carry this out, we created a model that could make micro-level predictions. We applied the Bass-ABM and traditional complex contagion models to identify those basic users who are more likely to adopt the premium membership after one month. In this comparison we did not include the complex contagion with innovation coefficient as there were no significant performance changes with respect to the traditional complex contagion model. However, we did build a simple logistic model with a ridge estimator (le Cessie and van Houwelingen 1992) using the Weka toolkit implementation (Witten and Frank 2005) and a random classifier as a baseline to help understand the performance of the agent-based DSS. The logistic model was designed to use one independent variable (number of premium friends of a user) and one dependent variable (if the user is adopting or not after one month). Again, a hold-out approach is followed by using 75% of the data for training and 25% for testing.

To validate the micro-level forecast performance of the agent-based DSS we defined a set of app users in the initial time period of one month and we label which of those users will eventually adopt. This dataset has 10,798 app users (basic users at the beginning of the month). We label those users who will adopt the premium membership after one month (718 from the total number of 10,798 users). The set of 718 premium adopters is called the true positives (TPs). The true negatives (TNs) are those who never convert in this month (the remaining 10,080 app users). We also calculated the false negative (FN) and false positive (FP) rates of the models. A FP occurs when a user is incorrectly predicted as a premium adopter when she actually never converts to premium. A FN is the opposite case where the user is incorrectly predicted to be a non-adopter when she does adopt. A good way to summarize this analysis is by using the receiver operating characteristic (ROC) curve and the area under the curve (AUC) (Witten and Frank 2005).

The ROC curve provides a way to represent the trade-off between FPs and TPs for different values of the rejection threshold by showing the relation between the sensitivity and specificity of the forecast. The AUC summarizes the area under the ROC in the entire range $[0, 1]$ of the FP rate. The higher the AUC value, the lower the FP rate for a given TP rate, i.e., the model performs better since it identifies true positives more frequently with less false positives. Figure 2 shows the ROC curves for Bass-ABM, complex, logistic, and random models. Additionally, its legend shows the AUC values. Complex contagion achieves a higher forecast performance than the Bass-ABM. In fact, the complex contagion model achieves better results than the logistic model. AUC values are also in line with the curves: complex contagion has the highest AUC value (0.7337) which surpasses both the Bass-ABM (0.601) and the logistic model (0.713).

## TARGETING AND INCENTIVE POLICIES USING THE DSS

Once the agent-based DSS is validated and modelers and marketers have accepted it, it is time to explore WOM strategies and scenarios of interest. This section tries to answer the main marketing questions of the app managers: how can we target and incentivize specific basic users to expand the premium market of the app? And, how many additional premium adopters can the *amplified WOM* of the incentivization policies encourage to convert? We use the agent-based DSS to select the best basic users to target by running simulations to evaluate the implications of different targeting policies for market expansion.

### *Using the DSS to Increase Social Influence via Rewards*

Before running the specific targeting policies, we use the DSS to understand how rewarding campaigns can increase the *amplified WOM* in the Animal Jam setting. This allows us to better understand how conversions lead to other conversions. The managers were interested in understanding how rewarding a user who converts to premium could affect the WOM effect on other

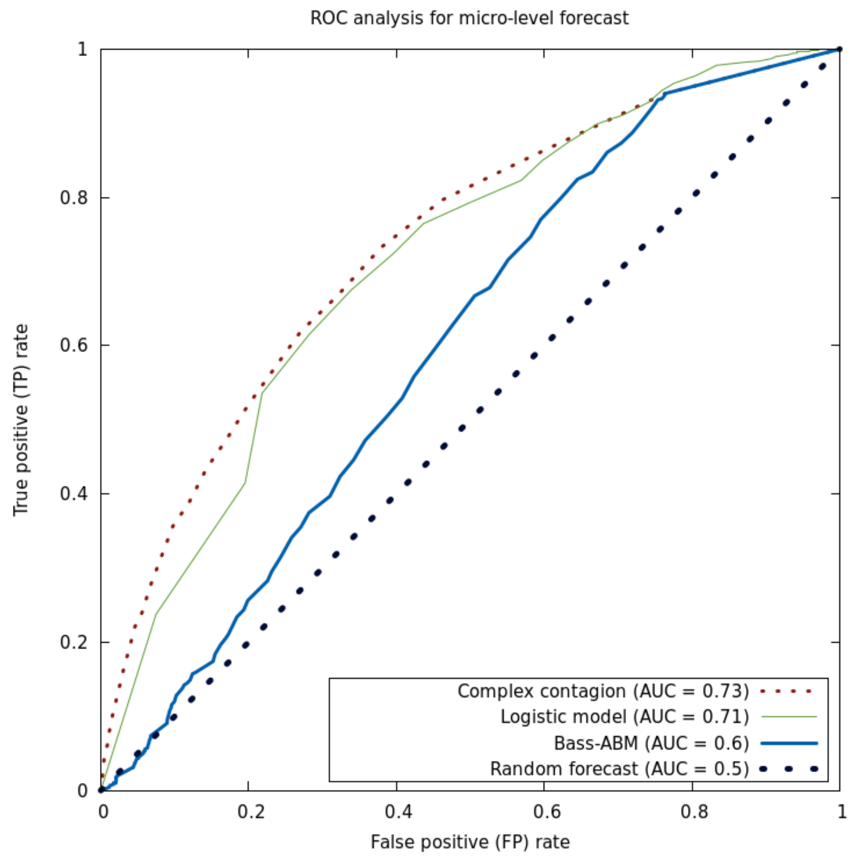


Figure 2: ROC analysis of the models for the micro-level forecast. AUC values are shown in the legend of the figure. ROC analysis of the models for micro-level forecast.

users potentially converting. After providing this reward, the agent-based DSS explores the social influence of the rewarded users with respect to their friends. The revenue of these policies is measured by counting the successful referrals after a given period of time. It is also important to notice that, in order to examine different types of rewards (i.e., app bonus, extra software features, or gifts), we explore different amounts of social influence by varying the social influence value per dollar value of the reward, i.e., we explore a variety of results based on how much a reward causes a user to talk positively about conversion. The idea being that if managers reward users when adopting, the rewarded users are more likely to talk positively about the benefits of a premium membership with their friends, increasing their influence over the baseline rate.

In order to accomplish our goal we used the calibrated Bass-ABM, i.e., the best diffusion model for fitting the macro-level data. Additionally, we made use of a weighted social influence (per dollar of investment) between the users to adjust the influence with their friends (parameter $\tau_{ij}$ defined in step $S2B$). These weights are included in the individual adoption probability rule of the Bass-ABM. This is done, for each basic agent, by multiplying the fraction of her/his premium contacts by the social influence weight of each contact. By using this method it is possible to amplify the influence of each converted premium user with their friends during the simulation time. At the beginning of the simulation, all the social weights $\tau_{ij}$ are set to 1 which means that initially all users affect each other uniformly. But when a basic user adopts premium content and is rewarded, her/his social influence with other users is increased (value greater than 1).

The app company does not empirically know the real social influence between two friends when rewarding one of them because they have never implemented that policy in the past. Thus to understand the potential range of effects, we ran a sensitivity analysis to understand the market expansion results when considering different values for the social influence weights. Marketers can look at these results, which include both pessimistic and optimistic scenarios, and make their decisions accordingly. If the app company implements this kind of rewarding policy in the near future it will be possible to analyze the app conversion data and estimate a value for the social influence weight per dollar spent on reward.

The plot of Figure 3 shows the number of additional premium users during 31 days in the test period by considering different increases in the social influence (x-axis of the plot) with respect to no rewarding (i.e., point $(0, 0)$). The main finding after analyzing this sensitivity analysis and scenario is that increasing social influence between users by rewarding users when adopting premium content has a positive non-linear impact on increasing the number of premium adopters of the app, i.e., referrals can be quite successful. We can also see how the lift in additional premium members does not have a linear behavior when increasing the social influence by rewarding them at the time of adoption.

*Targeting the Most Likely Users to Adopt Premium*

Furthermore, the Animal Jam managers wanted to explore if targeting basic users, as opposed to rewarding premium conversions, could also be a profitable strategy. Instead of rewarding users after they convert, we evaluated a marketing strategy to incentivize different groups of non-converted users (basic users) to stimulate them to adopt premium and to increase the *amplified* WOM of the app. We follow a similar targeting approach to the one of Haenlein and Libai (2013) by studying

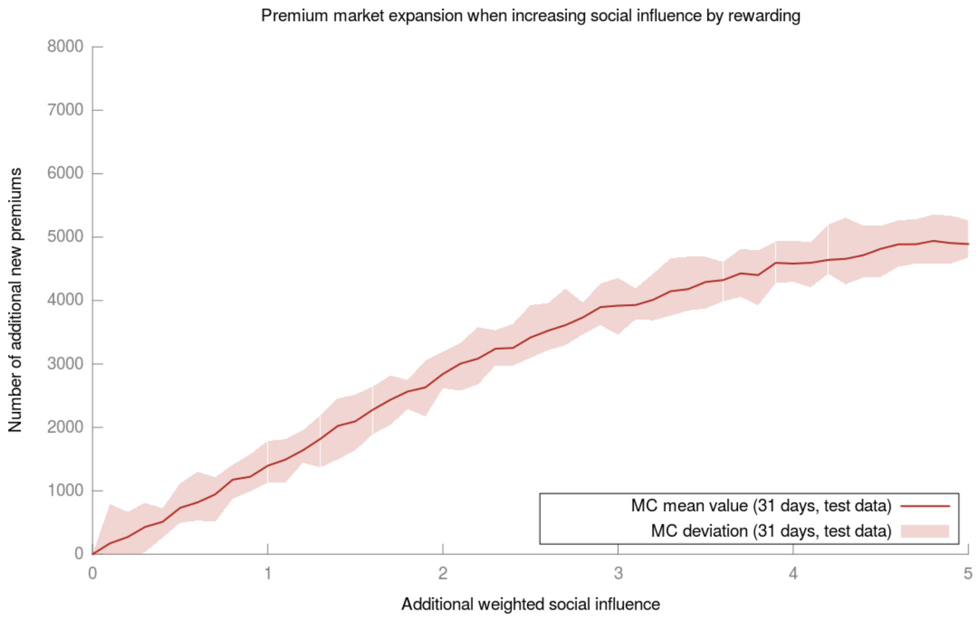


Figure 3: Agent-based DSS simulation output of additional premium adoptions when rewarding adopters versus the social influence per dollar (average and deviation of the 15 MC runs). The number of premium adoptions goes up if users exert more social influence per dollar but eventually it levels off.

the strategy to reward the most likely users to adopt, i.e., those expected to generate more value (*revenue leaders*).

To identify which users to target, we first examined how likely each user was to convert to premium on their own. As a result, we used the two-level forecasts of the agent-based DSS to first, target the users of the app and to second, evaluate the market expansion impact that happened as a result of their conversion by simulating the app market. The DSS generates a group of the 2,000 most likely basic users to adopt to be stimulated at the beginning. The user selection is done using the micro-level DSS forecast. As previously discussed, complex contagion is the best diffusion model for micro-level forecast and so we employ it for this analysis.

After targeting the group of basic users with incentives, we ran different simulations to forecast the additional premium conversions the company can obtain with respect to the case of not applying any incentive policy. A sensitivity analysis is run for the parameter that affected the weighted social influence ($\tau_{ij}$). Since the firm had never tried this before, we were not sure what the effect of providing an incentive would be on the innovation conversion rate. To solve this problem, we ran an additional sensitivity analysis on the lift in the innovation coefficient $\hat{p}$ when rewarding a target user.

Web Appendix D contains three heat-maps of the results after applying the targeting and incentive policies and running the DSS simulations. Each one is associated with a different way of creating the group of target users. The different groups are selected based on the complex contagion model parameter $\phi$ which discriminates whether a user is forecasted to adopt or not depending on the minimum number of adopters within her/his friends. Concretely, we examine 1, 3, and 6 for that threshold $\phi$ (minimum premium friends). Under some conditions (high social influence and high increase of the innovation coefficient), the incentivization policy allows the company to obtain more than $1,000$ additional premium conversions in one month.

For instance, when considering a scenario of a 0.5 lift in the innovation coefficient ($\hat{p}$) and 1.8 for the social influence weight ($\tau_{ij}$), we observe a considerable market expansion. In this case, the MC runs of the DSS result in an average additional premium conversions of 686.8 when targeting basic users with a minimum of 1 premium friends, in 713.33 when targeting basic users with a minimum of 3 premium friends, and 819.86 when targeting basic users with a minimum of 6 premium friends (see Web Appendix D). In general, the best expansion results are those when creating a target group of users with 6 or more premium friends, but the three scenarios present similar results.

*Managerial Implications of the DSS*

We created this agent-based DSS directly in response to questions that the app managers had about the best way to target and incentivize their basic users. Specifically we addressed the following questions that they had:

- *What is the importance of WOM and its role when free app users consider adopting premium content?* By analyzing the model's performance and sensitivity analysis on the diffusion parameters, the DSS showed that WOM influence has a significant effect on premium market expansion. For instance, we found, as a result of the complex contagion model, that the average user is likely to convert when three of their friends are premium.

- *What is the effect of rewarding conversions and how does this affect market expansion?* The exploration of different amounts of social influence, generated by rewards for conversion, showed that there is also a great potential to amplify market expansion when implementing such a strategy. However, we discovered that this exhibits is a non-linear relationship, and that the amount of expansion slows down as the value of influence per dollar goes up.

- *What are the benefits of running incentivization campaigns that target free users?* The agent-based DSS was able to evaluate the effect of different incentivization policies on basic users. An important insight for marketers was the need to target users with a high number of premium friends to generate more value through new premium adoptions: for instance, having six or more premium friends was the best tested policy.

## FINAL DISCUSSION

WOM programs and viral marketing have been studied in the past (Dellarocas 2006, De Bruyn and Lilien 2008, Van der Lans et al. 2010), and it has become clear that there are important concerns modelers must take into account when building a successful WOM-related DSS and to encourage a significant use by practitioners (Lilien 2011). Our methodological framework proposes an agent-based DSS to model the WOM dynamics and considers several guidelines that may allow the successful managerial adoption of their results. The cornerstone method of the framework, ABM, has already shown that it can capture both the social network structure and complex phenomena of customer interactions in previous research on WOM programs (Libai et al. 2013, Schlereth et al. 2013, Haenlein and Libai 2013), but in this paper we develop a generalizable way to build an agent-based DSS to assist WOM decisions and programs.

*Building and Using the DSS for a Freemium App*

We presented the application of the agent-based DSS framework to a real hedonic app, Animal Jam. The DSS application illustrated how to use the modeling guidelines and model construction steps to create a DSS for this particular app context, taking into account the data provided by the company. The Bass-ABM and two variants of the complex contagion were used to model the premium adoption of the app. The use of the complex contagion model (Centola and Macy 2007) came from the app data analysis since the adoption of premium content by a basic user appears to require a minimum number of premium contacts. We are not aware of any other study on marketing where adoption or WOM effects have been modeled using complex contagion, but it works well in this context and we encourage future researchers to explore the use of complex contagion models in other marketing research.

The use of a GA-based calibration method was helpful to validate the models of the DSS with respect to the empirical app data and to understand parameter variations and appropriate values of the models. The validation step showed that complex contagion achieves better results than the Bass-ABM and simple logistic models when forecasting if a user is going to adopt premium content in the near future (micro-level). Complex contagion supports the freemium model hypothesis of Bapna and Umyarov (2015) where the authors showed that the effect of peer influence is moderated by the number of friends of the user; but users with a smaller number of friends are more susceptible. However, complex contagion performs slightly worse for macro-level forecasting (fitting the historical premium evolution). Therefore, and although the differences in the macro-level forecasts are not huge, the best DSS configuration for the Animal Jam setting is to use the Bass-ABM to forecast the premium evolution of the market (macro) and the complex contagion model to individually forecast if a user is adopting (micro).

We examined a set of scenarios and policies by using the DSS to understand how to better expand the market by using different reward policies. We ran a sensitivity analysis to show that WOM leads to more premium members and enhances the effects of traditional activities (e.g., promotions or rewarding) as also demonstrated by Trusov et al. (2009), Wuyts et al. (2004). We confirmed that, for this freemium app setting, the lift in premium adoptions when rewarding users is not linear. This fact means when the social influence of users' relations overcomes a certain value, increasing the reward does not provide significantly increase premium expansion.

The joint use of micro- and macro-level forecasts in a validated DSS created a powerful decision-making tool to run marketing policies for a freemium business model. Both micro and macro-level forecasts are employed to simulate the effects of different targeting policies for the most likely basic users to adopt, a strategy to target revenue leaders (Haenlein and Libai 2013) (i.e., those expected to generate high profitability on their own). We compared three different configurations for targeting a group of basic users and we noticed a notable lift in the number of additional premiums if incentives affect social influence and innovation influence considerably.

*Limitations and Future Work*

Our hope is this study enhances research and development into agent-based DSSs and constitutes an important step toward the use of DSSs with high managerial success. In fact, there are many ways researchers can expand the capabilities of the framework by, for instance, including information about how *influential* a user is when targeting (Trusov et al. 2010). Following this research line, practitioners could compare the premium expansion implications of targeting users based on different social network features, similar to the high-degree and high-betweeness seeding strategies suggested by Hinz et al. (2011). We also assume that the adoption model of users is homogeneous across all users. Future work could explore the role of heterogeneity in adoption processes.

Although we built a DSS for a particular freemium app following the guidelines and proposed steps, the main purpose of this paper is to identify a set of guidelines and construction steps for the creation of a DSS for WOM programs. In future work, we would like to extend our work by considering more freemium scenarios and marketing insights such as the role of functionality and content updates, the customer perceived value of new software features, or the impact of including more social features versus non-social functionalities. We think a great deal of work is still needed to fully understand the best freemium model practices, but we believe that the DSS created here is a first step toward exploring these questions.

**References**

Anderson, C., 2009. Free: The future of a radical price. Hyperson: New York .

Axelrod, R.M., 1997. The complexity of cooperation: Agent-based models of competition and collaboration. Princeton University Press.

Bapna, R., Umyarov, A., 2015. Do your online friends make you pay? A randomized field experiment on peer influence in online social networks. Management Science 61, 1902–1920.

Barabási, A.L., Albert, R., 1999. Emergence of scaling in random networks. Science 286, 509–512.

Bass, F.M., 1969. A new product growth model for consumer durables. Management Science 36, 1057–1079.

Biyalogorsky, E., Gerstner, E., Libai, B., 2001. Customer referral management: Optimal reward programs. Marketing Science 20, 82–95.

Van den Bulte, C., Joshi, Y.V., 2007. New product diffusion with influentials and imitators. Marketing Science 26, 400–421.

Centola, D., Macy, M., 2007. Complex contagions and the weakness of long ties. American Journal of Sociology 113, 702–734.

le Cessie, S., van Houwelingen, J., 1992. Ridge estimators in logistic regression. Applied Statistics 41, 191–201.

Chai, T., Draxler, R.R., 2014. Root mean square error (RMSE) or mean absolute error (MAE)?–arguments against avoiding RMSE in the literature. Geoscientific Model Development 7, 1247–1250.

Cheng, H.K., Tang, Q.C., 2010. Free trial or no free trial: Optimal software product design with network effects. European Journal of Operational Research 205, 437–447.

Cheung, C.M., Thadani, D.R., 2012. The impact of electronic word-of-mouth communication: A literature analysis and integrative model. Decision Support Systems 54, 461–470.

Chica, M., Barranquero, J., Kajdanowicz, T., Cordón, O., Damas, S., 2016. Multimodal optimization: an effective framework for model calibration. Information Sciences .

Chica, M., Cordón, O., Damas, S., Bautista, J., 2010. Multiobjective, constructive heuristics for the 1/3 variant of the time and space assembly line balancing problem: ACO and random greedy search. Information Sciences 180, 3465–3487.

Chica, M., Cordón, O., Damas, S., Bautista, J., 2011. Including different kinds of preferences in a multi-objective ant algorithm for time and space assembly line balancing on different Nissan scenarios. Expert Systems with Applications 38, 709–720.

De Bruyn, A., Lilien, G.L., 2008. A multi-stage model of word-of-mouth influence through viral marketing. International Journal of Research in Marketing 25, 151–163.

Dellarocas, C., 2006. Strategic manipulation of internet opinion forums: Implications for consumers and firms. Management Science 52, 1577–1593.

Delre, S.A., Jager, W., Bijmolt, T.H., Janssen, M.A., 2010. Will it spread or not? The effects of social influences and network topology on innovation diffusion. Journal of Product Innovation Management 27, 267–282.

Dorigo, M., Maniezzo, V., Colorni, A., 1996. Ant system: optimization by a colony of cooperating agents. IEEE Transactions on Systems, Man and Cybernetics, Part B: Cybernetics 26, 29–41.

Epstein, J.M., 2006. Generative social science: Studies in agent-based computational modeling. Princeton University Press.

Feo, T.A., Resende, M.G.C., 1995. Greedy randomized adaptive search procedures. Journal of Global Optimization 6, 109–133.

Garcia, R., 2005. Uses of agent-based modeling in innovation/new product development research. Journal of Product Innovation Management 22, 380–398.

Godes, D., Mayzlin, D., 2009. Firm-created word-of-mouth communication: Evidence from a field test. Marketing Science 28, 721–739.

Goldberg, D.E., Holland, J.H., 1988. Genetic algorithms and machine learning. Machine Learning 3, 95–99.

Goldenberg, J., Han, S., Lehmann, D.R., Hong, J.W., 2009. The role of hubs in the adoption process. Journal of Marketing 73, 1–13.

Goldenberg, J., Libai, B., Muller, E., 2001. Talk of the network: A complex systems look at the underlying process of word-of-mouth. Marketing letters 12, 211–223.

Goldenberg, J., Libai, B., Muller, E., 2010. The chilling effects of network externalities. International Journal of Research in Marketing 27, 4–15.

Granovetter, M., 1978. Threshold models of collective behavior. American Journal of Sociology 83, 1420–1443.

Guidolin, M., Guseo, R., 2014. Modelling seasonality in innovation diffusion. Technological Forecasting and Social Change 86, 33–40.

Haenlein, M., Libai, B., 2013. Targeting revenue leaders for a new product. Journal of Marketing 77, 65–80.

Hinz, O., Skiera, B., Barrot, C., Becker, J.U., 2011. Seeding strategies for viral marketing: An empirical comparison. Journal of Marketing 75, 55–71.

Holland, J.H., 1975. Adaptation in Natural and Artificial Systems. The University of Michigan Press, Ann Arbor.

IGN, 2015. Blizzar gifts wow veterans with their own orc statue (last accessed 27-06-2016). Link: www.ign.com/articles/2015/01/19/blizzard-gifts-wow-veterans-with-their-own-orc-statue .

Kennedy, J., 2010. Particle swarm optimization, in: Encyclopedia of Machine Learning. Springer, pp. 760–766.

Kumar, V., 2014. Making freemium work: Many start-ups fail to recognize the challenges of this popular business model. Harvard Business Review 92, 27–29.

Van der Lans, R., Van Bruggen, G., Eliashberg, J., Wierenga, B., 2010. A viral branching model for predicting the spread of electronic word of mouth. Marketing Science 29, 348–365.

Lee, C., Gupta, V., Sunil, G., 2013. Designing freemium: a model of consumer usage, upgrade, and referral dynamics. Dissertation Harvard Business School (Job market paper) .

Leeflang, P.S., Wieringa, J.E., Bijmolt, T.H., Pauwels, K.H., 2015. Building models for markets, in: Modeling Markets. Springer, pp. 1–24.

Libai, B., Bolton, R., Bügel, M.S., De Ruyter, K., Götz, O., Risselada, H., Stephen, A.T., 2010. Customer-to-customer interactions: broadening the scope of word of mouth research. Journal of Service Research 13, 267–282.

Libai, B., Muller, E., Peres, R., 2013. Decomposing the value of word-of-mouth seeding programs: Acceleration versus expansion. Journal of Marketing Research 50, 161–176.

Lilien, G.L., 2011. Bridging the academic-practitioner divide in marketing decision models. Journal of Marketing 75, 196–210.

Lilien, G.L., Roberts, J.H., Shankar, V., 2013. Effective marketing science applications: Insights from the ISMS-MSI practice prize finalist papers and projects. Marketing Science 32, 229–245.

Little, J.D., 1970. Models and managers: The concept of a decision calculus. Management Science 16, B466–B485.

Lourenço, H.R., Martin, O.C., Stützle, T., 2003. Iterated local search, in: Handbook of metaheuristics. Springer, pp. 320–353.

Lovett, M.J., Peres, R., Shachar, R., 2013. On brands and word of mouth. Journal of Marketing Research 50, 427–444.

Luo, L., 2011. Product line design for consumer durables: an integrated marketing and engineering approach. Journal of Marketing Research 48, 128–139.

Macal, C.M., North, M.J., 2005. Tutorial on agent-based modeling and simulation, in: Proceedings of the 37th conference on Winter simulation, ACM. pp. 2–15.

Miller, J.H., 1998. Active nonlinear tests (ANTs) of complex simulation models. Management Science 44, 820–830.

Newman, M., Barabási, A.L., Watts, D.J., 2006. The structure and dynamics of networks. Princeton University Press.

Newman, M.E., 2003. The structure and function of complex networks. Society for Industrial and Applied Mathematics (SIAM) review 45, 167–256.

Oestreicher-Singer, G., Zalmanson, L., 2013. Content or community? A digital business strategy for content providers in the social age. Management Information Systems Quarterly 37, 591–616.

Oliva, R., 2003. Model calibration as a testing strategy for system dynamics models. European Journal of Operational Research 151, 552–568.

Peers, Y., Fok, D., Franses, P.H., 2012. Modeling seasonality in new product diffusion. Marketing Science 31, 351–364.

Rand, W., Rust, R.T., 2011. Agent-based modeling in marketing: Guidelines for rigor. International Journal of Research in Marketing 28, 181–193.

Rogers, E.M., 2003. Diffusion of innovations. 5th Edition. The Free Press. Simon and Schuster, New York.

Ryu, G., Feick, L., 2007. A penny for your thoughts: Referral reward programs and referral likelihood. Journal of Marketing 71, 84–94.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S., 2008. Global sensitivity analysis: the primer. John Wiley & Sons.

Sargent, R.G., 2005. Verification and validation of simulation models, in: Proceedings of the 37th conference on Winter simulation, pp. 130–143.

Schlereth, C., Barrot, C., Skiera, B., Takac, C., 2013. Optimal product-sampling strategies in social networks: How many and whom to target? International Journal of Electronic Commerce 18, 45–72.

Schmitt, P., Skiera, B., Van den Bulte, C., 2011. Referral programs and customer value. Journal of Marketing 75, 46–59.

Stonedahl, F., Rand, W., 2014. When does simulated data match real data? Comparing model calibration functions using genetic algorithms, in: Advances in Computational Social Science. Springer, Japan. volume 11 of *Agent-Based Social Systems*, pp. 297–313.

Stonedahl, F., Rand, W., Wilensky, U., 2010. Evolving viral marketing strategies, in: Proceedings of the 12th annual conference on Genetic and evolutionary computation, ACM. pp. 1195–1202.

Talbi, E.G., 2009. Metaheuristics: from design to implementation. John Wiley & Sons.

Teece, D.J., 2010. Business models, business strategy and innovation. Long range planning 43, 172–194.

Terano, T., 2008. Beyond the KISS principle for agent-based social simulation. Journal of Socio-informatics 1, 175–187.

Terblanche, N.S., 2015. Customers' perceived benefits of a frequent-flyer program. Journal of Travel & Tourism Marketing 32, 199–210.

Trusov, M., Bodapati, A.V., Bucklin, R.E., 2010. Determining influential users in internet social networks. Journal of Marketing Research 47, 643–658.

Trusov, M., Bucklin, R.E., Pauwels, K., 2009. Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site. Journal of Marketing 73, 90–102.

Trusov, M., Rand, W., Joshi, Y.V., 2013. Improving prelaunch diffusion forecasts: Using synthetic networks as simulated priors. Journal of Marketing Research 50, 675–690.

Venkatesan, R., Krishnan, T.V., Kumar, V., 2004. Evolutionary estimation of macro-level diffusion models using genetic algorithms: An alternative to nonlinear least squares. Marketing Science 23, 451–464.

Viger, F., Latapy, M., 2005. Efficient and simple generation of random simple connected graphs with prescribed degree sequence, in: Lecture Notes in Computer Science. Computing and Combinatorics. Springer. volume 3595, pp. 440–449.

Voinov, A., Bousquet, F., 2010. Modelling with stakeholders. Environmental Modelling & Software 25, 1268–1281.

Watts, D.J., Dodds, P.S., 2007. Influentials, networks, and public opinion formation. Journal of Consumer Research 34, 441–458.

Wierenga, B., van Bruggen, G.H., Althuizen, N.A., 2008. Advances in marketing management support systems. Springer.

Wilensky, U., Rand, W., 2015. Introduction to agent-based modeling: modeling natural, social, and engineered complex systems with NetLogo. MIT Press.

Witten, I.H., Frank, E., 2005. Data Mining: Practical machine learning tools and techniques (2nd edition). Morgan Kaufmann, San Francisco, USA.

Wuyts, S., Stremersch, S., Van den Bulte, C., Franses, P.H., 2004. Vertical marketing systems for complex products: A triadic perspective. Journal of Marketing Research 41, 479–487.

Xie, K.L., Xiong, L., Chen, C.C., Hu, C., 2015. Understanding active loyalty behavior in hotel reward programs through customers' switching costs and perceived program value. Journal of Travel & Tourism Marketing 32, 308–324.