The Effects of Word-of-Mouth: An Agent-Based Simulation of Interpersonal Influence in Social Networks

Dejan Duzevik, Andrey Anev, Pablo Funes, Paolo Gaudiano

August 21, 2007

Abstract

A recent article by Watts and Dodds raises the question of whether or not the spread of information and change of opinions in a social network depend on the existence of "influentials." We have developed a software simulation platform to analyze the impact of word-of-mouth and social networks on marketing and advertising. We use our simulation platform to test the relative impact of the structure of a social network and the nature of individuals within that network. Our results show that Watts' conclusions are only partly correct, and depend heavily on assumptions and on characteristics of the social network.

1 Introduction

There has been a growing interest among social scientists˜and particularly the marketing community˜in the notion of *influentials*, individuals who have a disproportionate influence on the spread of information and in the formation of opinions in their society [9, 7, 14, 10, 3]. For anyone interested in the spread of information among a human population, it is important to understand whether there are individuals who exert more impact on the knowledge, perceptions, and behavior than others. Social scientists interested in the causes of social change can find root causes in the existence, number, and behavior of influentials. The rapidly growing number of researchers interested in social networks can add another explanatory variable to the array of "physical" constraints of network structures. For marketers, politicians, and activists interested in actionable strategies, targeting influentials could increase the efficiency of their campaigns.

Despite the large interest and extensive research in this topic area, there are still fundamental questions that have not been answered: Do influentials exist, and what defines them? Do explanations of social phenomena such as diffusion of innovations, informational cascades, and public opinion formation need to assume the existence of influentials, or can those processes be modeled and explained without the use of influentials, as Duncan J. Watts and Peter S. Dodds (Watts and Dodds henceforth) have recently claimed [18]? If so, are there conditions under which influentials are more important than other variables that define a social network?

The difficulty in finding influentials is in part due to the different definitions. Are influentials opinion leaders, informed members of society, most connected people in a social network, trusted friends, celebrities, or some combination of these possibilities?

Another difficulty in identifying influentials is that, whichever way we define them, they have to be environment- and context-dependent: an IT professional may be well trusted and therefore an influential about computers in her network of friends, but may not have a large effect when it comes to health insurance; a highly connected lobbyist can change the minds of many politicians, but may not have any influence on the adoption of a new product. Therefore, because influentials can be many different things under different conditions, the questions about their existence and identification are best addressed through empirical approaches. Research on the population, on the process of information spreading, and on social networks is the most promising way to understand whether there are influentials and who they are.

Intuitively, the concept of "influential" seems sensible: we all know people who are more connected, convincing, trustworthy, informed, or simply more influential than others. However, the research needed to identify influentials is costly and time consuming. For this reason, before embarking on extensive research on influentials, it is important to understand whether the concept of influentials is actually useful in explaining, analyzing, and predicting any type of social change, especially in the context of changing consumer opinion about products. In the absence of any conclusive empirical evidence we need to understand under which conditions influentials matter.

In recent years computer simulation has emerged as a powerful approach to understanding and analyzing social networks. At Icosystem, we have applied a computer simulation technique known as *agent-based modeling* to the analysis of social networks for marketing and advertising. Although originally designed to help marketers with planning and strategy, our software platform is ideally suited to address research questions of the type raised by Watts and Dodds, because it captures in detail both the characteristics of individual consumers and the structure of the social network in which the consumers are embedded. In this paper we present a number of computer simulations that show the relative importance of variables related to the population, the network, and the process of influence itself on the dynamics of social networks.

Our results suggest that Watts and Dodds' results, as well as the outcomes from the influentials hypothesis, are only two of many possible outcomes that different mechanisms of interpersonal influence and network structure can have on the changing dynamics of a human population. More specifically, we show that in most scenarios we explore, the dynamics of opinion change is only mildly sensitive to the influence level of individuals. However, under certain conditions we find that when influentials contribute significantly to the dynamics of the system, the perception change in the total population is faster and greater than in scenarios where the influentials play a minor role.

Our results suggest that although labels such as "influential" may be useful in explaining concepts related to word-of-mouth marketing, the reality is that the success of marketing activities depends on a complex set of factors, and that focusing exclusively on targeting influentials while ignoring other aspects of social networks can be misleading. We believe that the most powerful approach is one that balances a mixture of market research, analytics, and computer simulations.

2 The Influentials Hypothesis and Watts and Dodds' Model

The *influentials hypothesis* states that "there is a minority of individuals who influence an exceptional number of their peers and are important to the formation of public opinion" [18]. The influentials hypothesis has its roots in two distinct research areas: the "new" science of networks [4, 2, 12, 17]and the slightly older theories of interpersonal influence [8, 15, 11, 13, 16]. The former builds on disparate fields such as discrete mathematics, physics, and sociology to understand the effects of network structure on the behavior of a population. The latter, mostly grounded in the social sciences, cognitive science and psychology, aims to uncover the mechanisms at play when people's perceptions and choices are influenced by others.

In models of social networks the structure of the network and the properties that define it are as important to the overall dynamics of the population as the properties of the individual people and the nature of their interactions [12]. Researchers in this area focus on the different dynamics originating from the network's clustering coefficients, node and edge proportions, edge distributions, and an array of variables that define the course of information-flow throughout the network.

The other precursor to the influentials hypothesis are theories of interpersonal influence that focus on identifying the processes and variables involved when two people exchange information and alter each other's knowledge, perceptions, or behavior. Seminal work on the effects of personal communication on people's preferences includes Paul Lazarsfeld and Elihu Katz's *Personal Influence* as well as Everett M. Rogers *Diffusion of Innovations* [13, 8]. The former argue that there exist scenarios where individual's opinions are influenced more by exposure to each other's opinions than by mass media messages, while the latter argues that the creation of the archetypal *S-curve* that represents a population's embrace of a novelty is caused by the existence of influentials.

Watts and Dodds' recent paper argues against the paradigm established by Lazarsfeld and Rogers and shows through a series of models that there are cases when the dynamics of opinion change can be explained by focusing on a "critical mass of easily influenced individuals," rather than using the influentials hypothesis. Watts and Dodds' model is a stride toward a more careful examination of the influentials hypothesis.

We believe, however that the conclusions from Watt's model are the result of some key assumptions. Specifically, Watts and Dodds make three important assumptions that set up his model. First, influentials are different from other members of the population only in the number of connections to others. Since the existing connections between simulated agents in Watts and Dodds' model determine whether influence will be exerted, the influentials of the model are the individuals with a larger number of opportunities to influence. What makes Watts and Dodds' definition of an influential difficult to apply to many real-world scenarios is the fact that even if we had a reliable way to observe the number of connections between the members of the population under study, we may still be wrong about how much they can influence the rest. Simply possessing an opportunity to influence may not necessarily translate into ability to affect the knowledge, perceptions, or behavior of those individuals.

Second, the process of public opinion formation in Watts and Dodds' model is a change in binary decisions that have positive externalities, *i.e.*, members of the population will switch their preference with a probability that is proportional to the number of neighboring members that share that preference. While it is true that this case is important and general enough to consider, we would like to explore processes that go beyond the case. For example, binary decisions do not cover environments of public opinion change when the opinions of a population are not yes/no choices but are better represented as a continuous variable (as is the case with the perception of competing brands or products). High involvement and high rationality categories of products (such as health insurance and financial services) are environments in which word-of-mouth and interpersonal influence are very important, yet consumer decisions are rarely based solely on what others around them are doing. More often than not, human decision-making is a complex heuristic that takes into account the decision-maker's perceptions about the alternatives. For this reason marketers and social scientists are often interested in tracking not solely the awareness of a certain alternative (a binary variable) but also people's perceptions and the intent to choose an alternative from the choice-space (both continuous variables).

Third, the network that underlies Watts and Dodds' model is a random network. Watts and Dodds test the effects of different variance around the mean number of edges, but the existing literature on network structure [1, 4, 2, 5, 19] gives us reason enough to believe that a different, non-Gaussian, distribution of edges may create altogether different dynamics. As described later, our simulations confirm this supposition.

3 Icosystem's Agent-Based Model

We have developed a software simulation platform to help marketers understand the impact of word-of-mouth, social networks and other types of consumer interactions. Our software is being used by three international advertising agencies and by a number of Fortune 500 corporations to plan and optimize advertising campaigns. A specific case study, which includes a description of the software tool, can be found in a companion paper [6].

Our software platform is based on the concept of *agent-based models* (ABM). In a nutshell, an ABM captures in software all the relevant elements of a system or process from the bottom-up. Each *agent* represents an entity in the system being simulated, including its main characteristics, its behaviors, and its responses to external factors. In most cases, and ABM includes multiple types of agents, representing all the relevant entities. Agents can include people (consumers, family members, friends, professionals, ...), companies (competitors, partners, vendors, ...), and other types of entities whose behavior can change during the course of a simulation (e.g., the media, manufacturers, retail points, ...). For this paper, we use the term "agent" to refer to a simulated consumer. The ABM simulates how all these elements and factors vary over time. At each time step, various events can take place: for instance, an agent may see an ad in a magazine, watch TV, talk with a friend, read an online review, or purchase a product. Each agent has a set of "preferences," which in this case can be thought of as a consumer's opinion about a product. A key aspect of our ABM is that the opinion of each agent fluctuates as a result of exposure to advertisements, communications with other agents, and even direct experience with the product.

For this paper we ignored advertising and direct experience with the product, and focused on the network structure and the mechanisms of personal influence. We start with the simple premise that agents have perceptions about the topic that is the subject of word-of-mouth communications. For simplicity, each agent's perception is represented by a value between 0 and 1. All agents have a number of social ties with others in the population. At each time step of the simulation, every agent has a certain probability of "talking" with another agent in its social network. In addition, each agent can belong to discernible groups, or segments. The notion of segments is useful because we can assign different levels of receptivity to influence depending on which segments the agents belong to. We run all simulations with two segments: *promoters* start the simulation with a perception of 1.0, while other agents start the simulation with a perception of 0.0.

Following Herbert Simon's formalism [16] that summarizes some of the early theories of social group interactions, we implemented the following mechanism of interpersonal influence:

$$p_A(t+1) = p_A(t) + [p_B(t) - p_A(t)]r_{AB}$$

where t is the current time step, p_A and p_B are represent the perception of agents belonging in segments A and B, respectively (for convenience, in what follows we sometime refer to "agent A" and "agent B" to mean "an agent in segment A" and "an agent in segment B," respectively) and r_{AB} is agent A's receptivity to B's influence, or in other words B's influence over A. When agent A communicates her perception to agent B, agent B's perception will change in the direction of A's perception, by an amount that depends jointly on the discrepancy between their opinions and on the receptivity of segment A toward segment B. Hence, if agent A and agent B have the same perception, the effect that their communication has on their perceptions is zero. When agent A is receptive to the opinion of agent B ($r_{AB} > 0$) and A's perception is different from B's perception, after a communication A's perception will move closer to B's perception the more receptive A is and the greater their opinion discrepancy.

Notice that by using the concept of segments we can allow for the asymmetry of communications that occurs with influentials: for example, the receptivity r_{AB} of agent B (the influential) to information from agent A (the influenced), can be much weaker than the receptivity r_{AB} of agent B to agent A. In fact, for the purposes of this paper we can set $r_{AB} = 0$, meaning that the influentials are not influenced at all by communicating with other agents.

Another important distinction is our use of the terms *promoter* and *influential*. By "promoter," we refer to an agent whose perception of a given product is very high (p is close to 1). By "influential," we refer to an agent whose opinion has a high influence on other agents (r close to 1). Presumably, in real life some people are "hard-wired" as influentials, while anyone can become (or stop being) a promoter as a result of external factors such as experience with the product or information exchange.

3.1 The processes of information exchange

Our model of perception change captured in the equation above, produces diminishing returns to each additional exchange of information, as shown in Fig.1. The concept of diminishing returns, which is fundamental to economic theories, is closely related to the concept of "response curves" used by marketers and media planners focused on optimizing advertising spend across different media channels. In that context, response curves represent the additional impact gained by additional spend during a marketing campaign. In this paper, the influence is the result of communications between agents rather than exposure to advertising. However, our platform is able to capture both types of response curves: in effect, communication between consumers can be thought of as a different channel with different response curve characteristics.

Two indicators are crucial to define a response curve: *potential* and *speed*. The potential of the curve is the maximum impact that a marketing campaign, or in our case the word-of-mouth interactions, can have on the average perception of the population. It is the point on the curve where no more uplift is achieved with additional spend, or in our case, additional exchange of information. The speed of the curve indicates the time that it takes to reach the potential or some portion of the potential (in the following simulations we define the speed of a curve as the time it takes to reach 90% of the potential). The speed parameter determines the shape (steepness) of the response curve, and essentially represents how quickly information changes the perception of an agent. In our model, the speed is given by the receptivity parameter r.

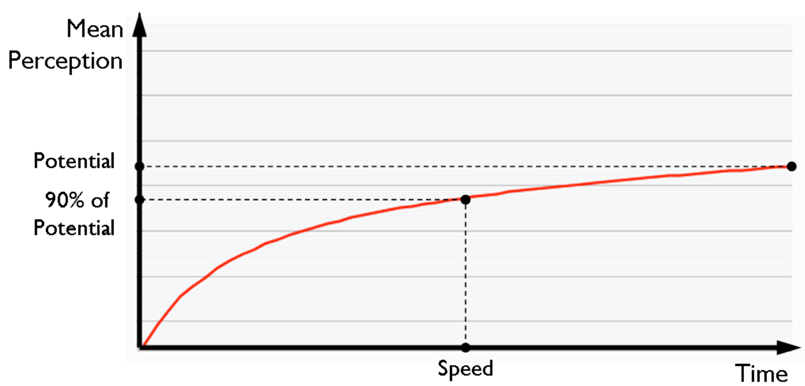

Figure 1: The perception of an agent resulting from exposure to a higher perception.

In an earlier section we mentioned Rogers' notion of the S-curve as representing the diffusion of innovations through information spread. The process of information spread

(S-curve) and the perception change (diminishing returns, or response curve) are not mutually exclusive: the response curve represents each agent's change in perception as a result of its exposure to information; while the S-curve is an *emergent property* of the entire social network, and it measures how information is spreading through the network.

To see both curves, we model a process where a minority of promoters (which, as you recall, are individuals who have a high perception) are introduced within a much larger population of individuals who are not aware of the product and have a very low opinion of it once they are informed of its existence. Figure 2 shows two such curves obtained during a typical simulation: the curve labeled "perception" shows the average perception of agents across the population (response curve), while the curve labeled "awareness" shows the spread of information through the population (in terms of the percentage of the agents who have been exposed to the information).

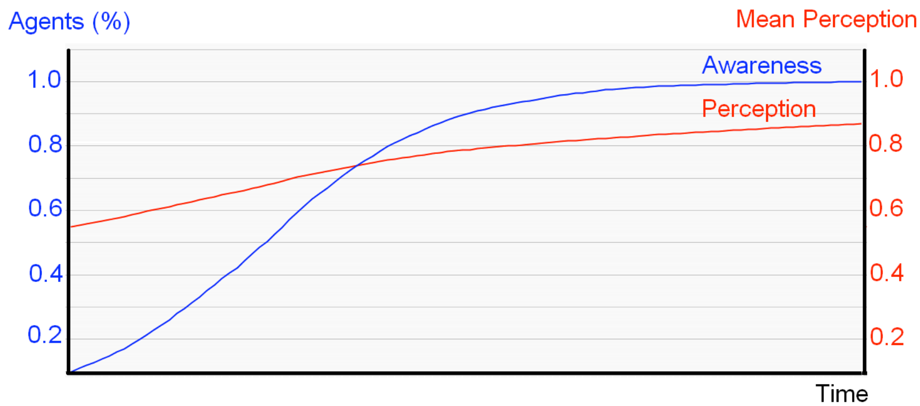

Figure 2: S-Curve and Response Curve

Having explained our interpretation of information exchange, we can start examining some of the outstanding questions regarding the "influentials hypothesis."

3.2 Influentials Definition and Network Type

Watts and Dodds claim that public opinion formation is driven by a critical mass of easily influenced individuals and not by influentials. Yet, many easily influenced individuals or a few very influential individuals are simply two sides of the same coin. A more interesting question is whether there are conditions under which the change in public opinion is a result of a group of individuals who are not necessarily much more influential than their peers. Furthermore, are there conditions under which the change in public opinion can only be driven by individuals who exert a disproportionate effect on the population? Finally, is there a trade-off between the number of non-influential promoters and the influence level of influentials?

In our simulations, we are interested in measuring how the average perception across the entire population changes as a function of the number of promoters. However, we are interested in investigating how perception is also influenced by two distinct factors. First, to test Watts' claims, in one set of simulations we assume that the level of influence of an agent is solely due to the number of connections with other agents. In another set of simulations, the level of influence is determined by the extent to which the agent can bring other agents' perceptions closer to her own. Second, we are interested in the difference that the underlying network structure may have on the dynamics of the system. As an initial step, we have simulated a random small-world network (similar to Watts' network) and a scale-free network. Before providing details of the simulation setup, we describe briefly these network types.

Figure 3 shows histograms of the number of connections per agent for the two different network types. Each bar in the histogram represents how many agents have that number of connections. The left histogram, representing a random small-world network, shows a bell-shaped distribution, with all agents having between 1 and 8 connections; the right histogram, representing a scale-free network, shows a more typical power-law distribution, in which most agents only have one or two connections, but some agents can have many connections — in this case as many as 494 (with a population of 1,000 agents).

The shading in the histograms of Fig.3 show how the population is divided into quintiles, *i.e.*, it shows that the first 20% of the agents in a random network have mostly 1, 2 and sometimes 3 connections; the second 20% have sometimes 2 and sometimes connections; and so on up to the top quintile, in which agents can have 6-8 connections. In a scale-free network, in contrast, the first two quintiles and part of the third quintile only have one connection — in other words, nearly 50% of the population. This type of distribution is actually typical of several types of social networks.

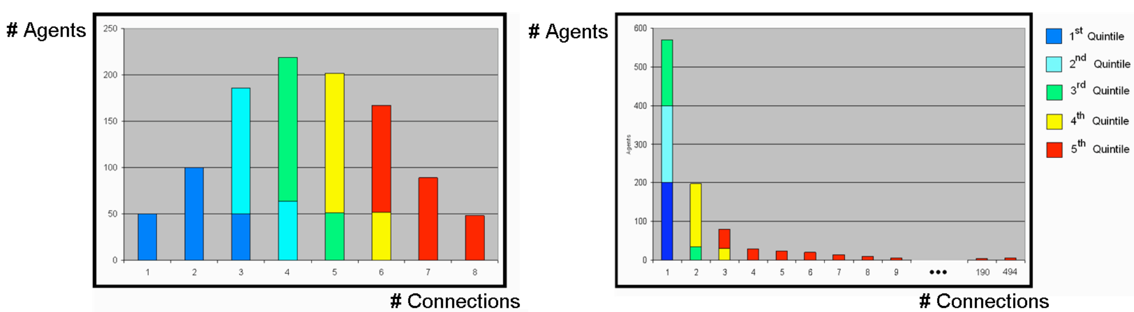

Figure 3: Random (small-world) and scale-free network distributions.

Figure 4 shows a diagram of the 2-by-2 sets of simulations we executed. For each set of simulations we generate a matrix of points 25 points. Each point represents the overall perception across the population of agents. In all four simulation sets, the Y axis measures the effect of changing the number of promoters (3, 6, 9, 12 or 15% of the population). In two of the simulation sets (top row of Fig.4), the X axis represents the number of connections that the promoters have. Specifically, we choose one of the five quintiles shown in Fig.3, and assume that all promoters belong in that quintile. Hence there are five points along the X axis, representing an increasing number of connections for the promoters. In the other two sets (bottom row of Fig.4) we assume that the promoters are distributed randomly in the network, but we change the amount of their influence on other agents (r in the equation) for five distinct values: 0.1, 0.3, 0.5, 0.7, and 0.9. Hence each set of simulations provides a 5x5 grid of points.

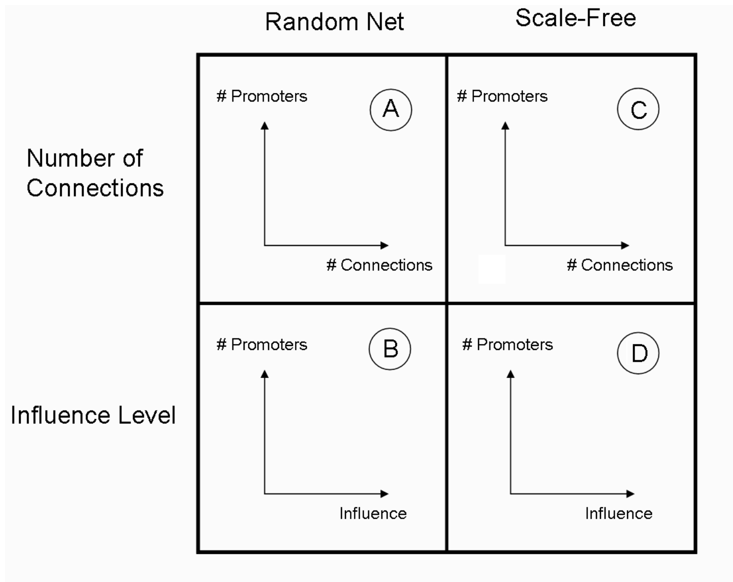

Figure 4: Our 2-by-2 simulation design.

Furthermore, the value of each point in each simulation set is actually the average of a Monte Carlo simulation (multiple runs with different random seeds) to minimize noise and the effect of chance.

For all simulation we recorded the speed and the potential of the change in average perception. Again, greater potential means that the final perception across the population is higher (in other words, the word-of-mouth was more effective in convincing people), while greater speed means that the change in perception spread through the population more rapidly.

In essence, the results of the experiment represent a type of sensitivity analysis that shows how the number, the connectivity, and the level of influence affect the population's change in perception for two network types and two processes of influence (the influentials hypothesis vs. Watts' hypothesis).

Before proceeding to the results, a note on the plots. The effects on perception are shown in all subsequent figures as contour plots. Each point on a plot is a simulation run. The colors show the change in perception. The two extremes on the spectrum are violet and red signifying low perception change and high perception change, respectively. If the two variables were equal in their effect on perception change the lines delimiting the color regions would be straight and run diagonally from the top left to the bottom right; if one variable were an absolute driver of the system's dynamics then the lines would run perpendicular to the axis of the variable that is dominant (in other words, there would only be changes in perception as a function of changes in that one variable).

3.3 Simulation Set A: number of promoters vs. promoter connectivity in a random network

For this simulation set we replicate the assumptions from Watts' model. First, we define influence as a property that is proportional to the number of connections in the social network. Second, the network is random with the mean number of edges per node $k = 3$. We assume that 10% of the population is promoters (their perception is 1.0 at the beginning of the simulation) and the rest of the population has perception of 0.0. We maintain a constant probability that two agents will talk to each other during a given time step at 10% throughout all the simulation sets. The probability that two agents will talk is calculated independently for each agent connection, so that agents with more connections will have more opportunities to influence the rest. The variables whose effects we aim to explore are the number of promoters (ranging from 3 to 15 percent of the population) and the quintile from which their number of connections is drawn. When all the promoters are selected from the highest quintile, they all have a large number of connections, and therefore they are "influentials" according to Watts' definition. Hence moving along the X axis corresponds to having more influential promoters.

In the context of Watts' assumptions (random networks, and influence being proportional to the number of connections), if the influentials hypothesis is correct, then the variation in outcomes of different runs should be more sensitive to the number of connections than on the number of promoters. Hence the contour plots should show lines that have a steep slope with respect to the x-axis (or be parallel to the y-axis, if connectivity drives the dynamics completely). If Watts' claims are correct, then we should see the opposite: the system should be more dependent on the number of promoters than on their degree of connectivity, and the contour plots should have a more gentle slope from left to right (or be parallel to the x-axis if the number of promoters was completely responsible for the change in effects).

The results for simulation set A (Fig. 5) support Watts' hypothesis more than they do the influentials hypothesis, but not entirely. The number of promoters accounts for a significant part of the perception change. For low proportions of promoters (3% to 9%, lower end of the Y axis), increasing the connectivity of the promoters (left to right on the X axis) only leads to a small change in the perception. However, as the number of promoters increases, the slope of the contour plot increases noticeably, meaning that the greater the number of influencers, the greater the final perception. In fact, one can see from the plot in Fig.5 that the highest average perception that a fifteen percent of not-well-connected promoters can achieve is around 0.31 (or 31 percent increase in perception) while fifteen percent of well-connected promoters can achieve 0.47 perception change. Moving the promoters from the lowest connected quintile to the highest boosts the promoters' effectiveness by 16%.

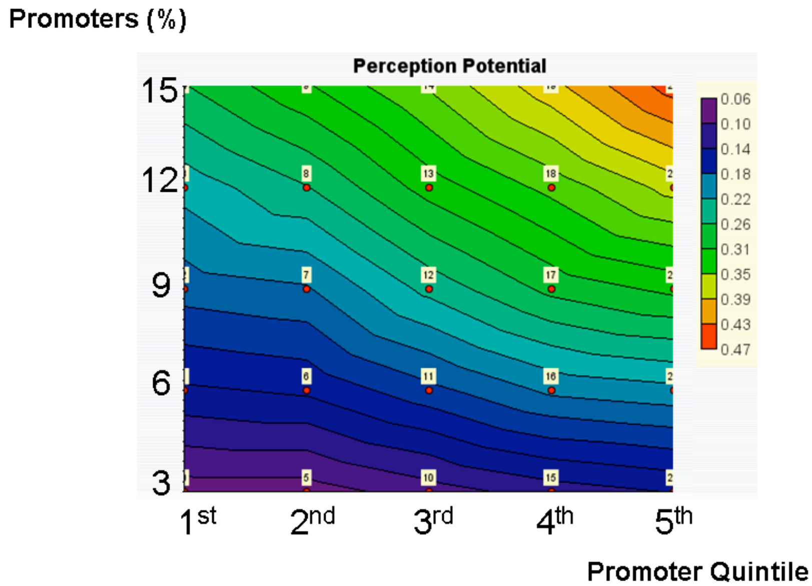

Figure 5: Number of Promoters Vs. Connectivity in a random network.

In conclusion, when promoters are few in number, the extra influence (connections) does not have a large effect on the outcomes. However, when promoters are more numerous, the extra influence can significantly amplify their effect.

3.4 Simulation Set B: number of promoters vs. promoter influence in a random network

In this simulation set we relax Watts' assumption that influence comes solely from the number of connections. We assume instead that when agents communicate some of them can exert greater influence than others - regardless of the number of connections they have. This corresponds to changing the variable r in Equation 1.

The promoter's influence (or the others' receptivity to promoters' opinion) is one of the experimental variables, with a range of 0.1 to 0.9. If an agent has influence of 1 over another, it means that when that agent talks to another, the "listener" will switch perception instantaneously to match the "speaker." Conversely, zero influence means that the agent will never be able to change another agent's perception. The values between the extremes represent the rate with which a perception change will occur. From Equation 1, if an agent with influence 0.5 and perception 1.0 "talks" to an agent with perception 0.0 five times, at the end of the fifth time-step the latter agent's perception will be 0.9325 if no other conversations took place.

At the start of each simulation, the influence of all agents over each other is set to 0.1, while the promoters' influence varies in the simulations. Both promoters and other agents in the simulation have the same probability to be well-connected or not (*i.e.*, they are drawn equally from all connectivity quintiles). Therefore, the only variable that makes some agents more influential than others is the level of influence r.

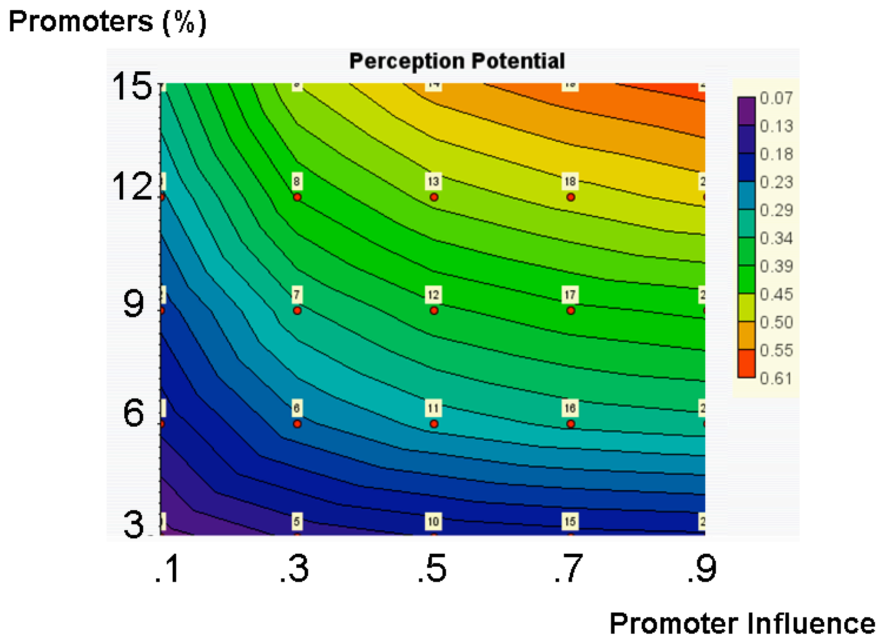

Figure 6: Number of promoters Vs. influence in a random network.

With our new definition of "influential," we can test the influentials hypothesis again. Figure 6 again shows the results of 25 simulation runs. Watts' claim that the number of promoters is more important than the level of influence is again partly corroborated by the results. However, when the influence level is below 0.3, the level of influence of promoters has a significant impact on the results, as evidenced by the contour lines having a steep downward slope. The influence level of 0.3 is a sort of tipping point, after which the number of promoters becomes the dominant driver of effects (shallower slope). Similarly to the results with Simulation Set A, however, the level of influence becomes more important for greater numbers of promoters.

In sum, the change of definition of an influential provides some support for the influentials hypothesis, but only under a restricted range of parameters. When promoters are few, extra influence does not have a large effect on the change in perception. When they are more numerous, like in Simulations set A, the extra influence adds to the overall effect.

3.5 Simulation Set C: number of promoters vs. promoter connectivity in a scale-free network

Simulation Set C is identical to Simulation Set A except that the underlying network is a scale-free network instead of a random network. As shown earlier in Fig.3, the edge distributions between the two types of networks are very different. In a random network, the most connected agent has 7 connections, while in a scale-free network of the same size the most connected agent has 493 connections. Numerous studies have shown that for processes that are binary (such as spread of viruses, raising awareness, and spread of innovations) the effects on the system are highly sensitive to the most connected nodes (or hubs).

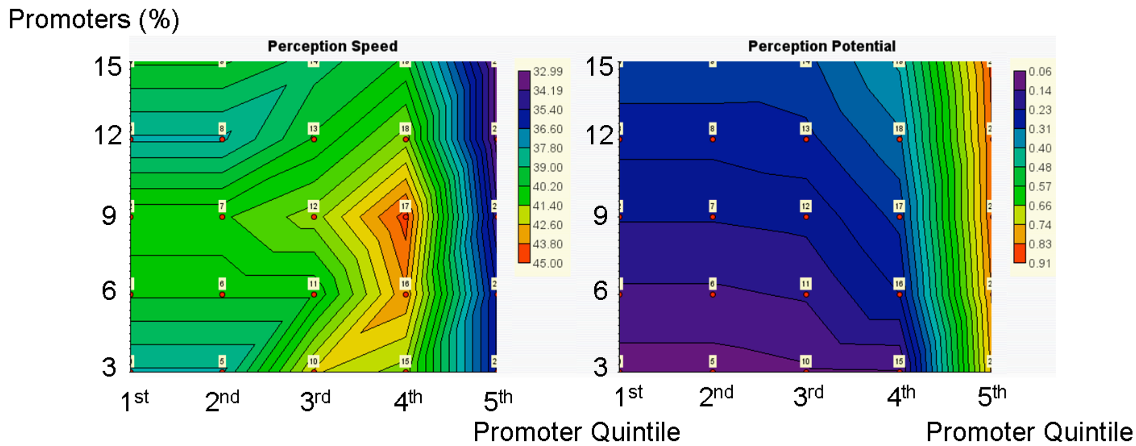

Figure 7: Number of promoters Vs. connectivity in a scale-free network.

Fig.7 shows that the system was insensitive to the number of promoters[1]. The sole driver of influence was the connectivity of promoters. In addition to a contour plot showing the effects on perception potential, we also included the perception speed, which follows the same pattern. The results of this simulation set are not surprising, but they are a reminder that changing a key assumption may produce a completely different set of outcomes. Whether the influentials hypothesis applies or not is highly dependent on the structure of the underlying network. The results are not inconsequential, given that scale-free networks are ubiquitous among online sites and in systems where some individuals have disproportionately frequent contact with others. An example of the latter are doctors as part of a patient's network of advisees about medical insurance or different prescription drugs.

3.6 Simulation Set D: number of promoters vs. promoter influence in a scale-free network

The promoter's influence is certainly important when it comes to the ability to alter the perceptions of a population when the underlying network is a scale-free. Does this outcome hold, when influence is defined as the ability to change others' perceptions and promoters and other agents are randomly assigned connections from the edge distribution?

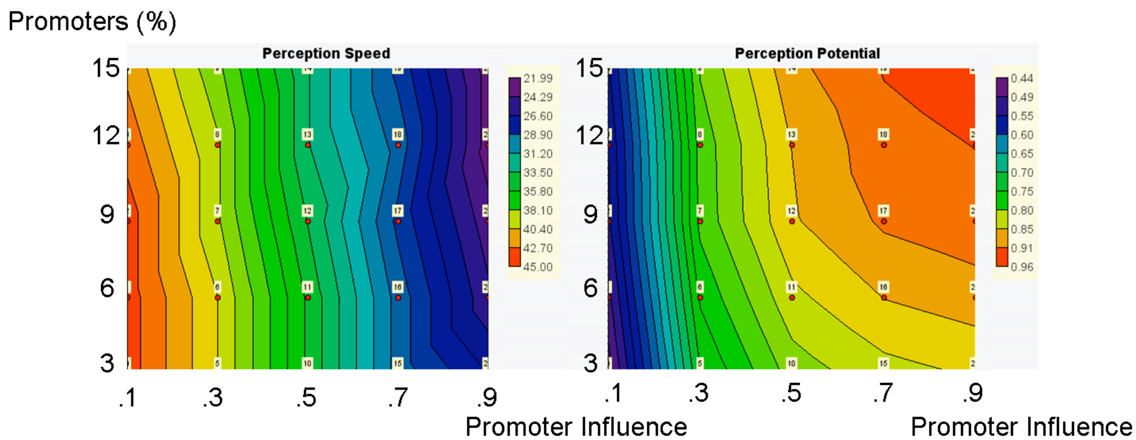

Figure 8: Number of promoters Vs. influence in a scale-free network.

Figure 8 shows a trade-off (in terms of perception potential) between number of promoters and level of influence almost identical to that observed in the corresponding simulation set with a random network. Fifteen percent of not very influential promoters can achieve the same perception uplift as six percent of influentials. The influentials hypothesis does not hold because the dynamics of the system can occur without the presence of any members who have disproportionate influence. At the same time, when individuals with greater influence exist in the system, their effects on the system's dynamics are not inconsequential.

3.7 Summary of simulation results

Three simulation sets (A, B, and D) showed that most outcomes in our system can be produced without the existence of influentials, but that when promoters were also more influential than other agents they reached the highest potentials. Simulation set C fully supported the influentials hypothesis and demonstrated that when promoters are disproportionately more influential than others (as the scale-free network enabled them to be) the system is highly sensitive to the influence level. Interestingly, the simulations in set C reached the highest uplift in perception of all runs (Table 1). The best run was the one in which 15% promoters from the most-connected quintile moved the average perception from 0.15 to 0.91. What this suggests is that hubs in a scale-free network are the most efficient way to increase the average perception in a population. While the influentials hypothesis does not apply in many environments, when it does apply its effect is much larger than the effects of the rest of the dynamics in the system.

Simulation Run	Maximum Potential
A. Promoter Number Vs. Connectivity (Random Net)	.47
B. Promoter Number Vs. Influence (Random Net)	.61
C. Promoter Number Vs. Influence (Scale-Free Net)	.91
D. Promoter Number Vs. Connectivity (Scale-Free Net)	.51

4 Conclusions

In this paper we have used a simulation tool to provide some quantitative insights into the influentials hypothesis, and Watts' recent claim that influentials are not as relevant as commonly thought.

Our simulations lend some support to Watt's claims, but more importantly they lead to some important observations.

First, social networks exhibit rich dynamics that depend in a complex way on a variety of factors. Any theory, hypothesis or claim that is based on "categorical" notions is likely to be at least partly misleading. We have shown for example that Watts' claims are largely correct in some regions of parameter space, but that his claims break in some portions of the parameter space.

Second, the structure of the underlying network is extremely important — possibly more so than the definition or even the magnitude of the influence between individuals. This was most clearly evident in simulation set C, in which the joint assumption of influence being proportional to connections, and a scale-free connectivity pattern led to results in support of the influentials hypothesis. It can be argued that this is an artifact resulting from the assumption that the probability of communications is independent for each connection (which means that the most heavily connected agent on average has nearly 500 times as many conversations as the least connected). However, this simply strengthens our previous comment that categorical statements and hypothesis should be regarded with some skepticism.

The most important conclusion in our opinion is that marketers who wish to use word-of-mouth as a means of influencing consumers should be careful to understand as much as they can about their consumers. They should also be sure that they understand the implications of their assumptions.

Finally, we hope that this paper demonstrates the importance of a quantitative approach to addressing some of these issues. More importantly, we believe that the agent-based formalism is ideally suited to addressing these and other challenging issues with which marketers are faced today. Traditional statistical approaches are too limited by their assumptions of linearity and stationarity to be of any substantial value in unraveling the complexities of today's consumer society.

List of Figures

References

[1] Reka Albert, Jeong Hawoong, and Albert-Laszlo Barabasi. Internet: Diameter of the world-wide web. *Nature*, 401:130 — 131, 1999.

[2] Barabasi Albert-Laszlo. *Linked*. Penguin, New York, NY, 2003.

[3] David Balter and John Butman. *Grapevine: The New Art of Word-of-Mouth Marketing*. Penguin Group, New York, 2005.

[4] Albert-Laszlo Barabasi and Eric Bonabeau. Scale-free networks, 2003.

[5] Tian Bu and Don Towsley. On distinguishing between internet power law topology generators. IEEE, 2002.

[6] Paolo Gaudiano, Oliver Bandte, Dejan Duzevik and Andrey Anev. How word-of-mouth impacts medicare product launch and product design. Word of Mouth Research Symposium 2007.

[7] Malcolm Gladwell. *The Tipping Point: How Little Things Can Make a Big Difference*. Little, Brown, and Company, New York, NY, 2000.

[8] Elihu Katz and Paul F. Lazarsfeld. *Personal Influence*. The Free Press, Glencoe, IL, 1955.

[9] Ed Keller and Jon Berry. *The Influentials*. The Free Press, New York, NY, 2003.

[10] Justin Kirby and Paul Mardsen. *Connected Marketing: The Viral, Buzz and Word of Mouth Revolution*. Elsevier, Burlington, MA, 2006.

[11] Paul F. Lazarsfeld, Berard Berelson, and Hazel Gaudet. *The People's Choice: How the Voter Makes up His Mind in a Presidential Campaign*. Columbia University Press, New York, NY, 1968.

[12] Mark Newman, Albert-Laszlo Barabasi, and Duncan Watts. *The Structure and Dynamics of Networks*. Princeton University Press, New Jersey, 2006.

[13] Everett Rogers. *Diffusion of Innovations*. Free Press, New York, NY, 1962.

[14] Emanuel Rosen. *The Anatomy of Buzz: How to Create Word-of-Mouth Marketing*. Doubleday, New York, NY, 2000.

[15] John Scott. *Social Network Analysis: A Handbook*. Sage Publications, London, UK, 1991.

[16] Herbert A. Simon. *Models of Man*. Wiley, New York, NY, 1957.

[17] Duncan J. Watts. The "new" science of networks. *Annual Review of Sociology*, 30:243—270, 2004.

[18] Duncan J. Watts and Peter S. Dodds. Influentials, networks, and public opinion formation. In Press, 2007.

[19] Huberman Bernardo A. Adamic Lada A. Wu, Fang and Joshua R. Tyler. Information flow in social groups. *Physica A Statistical Mechanics and its Applications*, 2004.

[1] *We also tested whether these finding change for different receptivity combinations and found them to hold*