CHAPTER 7

METRICS FOR THE NEW INTERNET MARKETING COMMUNICATIONS MIX

RANDOLPH E. BUCKLIN, OLIVER J. RUTZ, AND MICHAEL TRUSOV

Abstract

The Internet provides marketers with an expanded set of communications vehicles for reaching customers. Two of the most important and fast-growing elements of this new communications mix are online advertising and electronic word-of-mouth. While these vehicles provide new opportunities, they are also challenging marketers to understand how consumers respond and to develop new metrics for assessing performance. The purpose of this chapter is to review recent research developments in marketing that are most relevant to assessing the impact of these communications vehicles.

The chapter first discusses the two major forms of Internet advertising, display ads (also known as banners) and paid search. (Paid search refers to the text ads served as sponsored links by Internet search engines.) The existing literature on banner ads provides a body of empirical findings as well as a set of methods that marketers can draw upon to assess performance. Research on paid search is still emerging but early work has developed approaches that marketers can use as the basis for suitable performance metrics.

As a social medium, the Internet offers users new ways to communicate more easily and more extensively with others. Online communities, social networking sites, online referral programs, product reviews, and blogs all allow word-of-mouth to spread faster and farther than in the past. Research has shown how electronic records of online word-of-mouth (e.g., product reviews) can be connected, via models, to performance outcome variables such as product ratings and sales levels. Emerging work shows promise at identifying the particular members of online communities who are most likely to influence the actions of others, revealing opportunities for firms to manage these interactions.

Introduction

New ways of reaching customers via the Internet enable marketers to move beyond the traditional marketing communications mix. Two of the most important and fastest growing of these new communications vehicles are Internet advertising and the Internet as a social medium for transmission of word-of-mouth. These new communications vehicles pose challenges to marketers in understanding consumer response and in developing performance metrics. The purpose of this chapter is to examine some of the existing and emerging findings and approaches in this domain. In so doing, we seek to aid both academics and practitioners as they begin to address the challenges posed by the new Internet communications mix.

In the first part of the chapter, we will cover developments in Internet advertising. We discuss consumer response to and metrics for assessing the two major elements of paid Internet advertising: (1) display (or so-called "banner") ads, and (2) paid search advertising. Paid search refers to the text ads served as sponsored links in response to a user's inquiry at search engines such as Google and Yahoo! These are presented alongside nonpaid listings (known as organic or natural search results), which are based on the search engines' algorithms.

The second portion of the chapter will cover Internet word-of-mouth, social networks, and user-generated content. For example, marketers now have the opportunity to track various forms of word-of-mouth (WOM) communications over the Internet (e.g., blogs and electronic referrals). We will examine metrics for WOM on the Internet and the role that social networking sites (e.g., Friendster, Facebook) and user-generated content are beginning to play.

Understanding the effects of both advertising and WOM (and the proper metrics to assess them) were important topics long before the advent of the Internet. Thus, it is worth asking, precisely what is new due to the Internet medium? Our view is that two factors produce a new set of challenges. The first factor is that the medium itself is different. Almost by definition, consumers therefore respond to it (and interact with it) in novel ways. For example, studies have already documented how many consumers use the Internet in a goal-directed fashion (e.g., Moe 2003; Montgomery et al. 2004), suggesting that it is a potentially more involving medium than, say, television or print. The second factor is that the Internet allows consumer activity to be examined at a granular level. So-called clickstream data—the electronic trace of what users do and are exposed to at a website—allow marketers to consider the effectiveness of advertising and/or the impact of WOM at not only the individual level but at the level of the individual click. While this sets up unprecedented targeting opportunities, it also creates challenges for reliably measuring the effects of marketing communications at such a micro level.

Before proceeding, we note that Internet advertising and online WOM are just two elements of new Internet marketing communications mix. Broadly speaking, any Internet-based communication instigated by marketers or with marketing implications could be considered part of this mix. For example, companies communicate with their customers via their own websites and by hiring representatives to chat online with customers. E-mail campaigns organized to achieve customer acquisition and retention are now widespread and becoming more sophisticated. In this review, we focus on advertising and WOM. In both cases, relative managerial importance is high (i.e., spending levels and growth) and there is a critical mass of research findings and methods that are relevant for performance assessment and metrics development.

Internet Advertising

Internet advertising has now established itself as an important and fast-growing part of the marketing mix. Spending in the United States on Internet advertising increased from $4.6 billion in 1999 to $16.9 billion in 2006 (PricewaterhouseCoopers 2006) and the trend continues strongly upward. For example, the first quarter of 2007 saw a 29 percent increase over the first quarter of 2006. This spending on Internet advertising is composed of two major categories: display advertising and search engine advertising (paid search). In Internet display advertising consumers are exposed to ads that appear as banners on a web page or in pop-up ads that open in a separate browser window. On the other hand, paid search is a service offered by Internet search engines (e.g., Google and Yahoo!) in which the advertiser selects specific search terms (known as keywords) and creates a text ad designed to be seen by users who submit search queries containing the terms.

Display advertising had accounted for most online advertising spending until 2005, when the faster-growing paid search advertising overtook it. Interestingly, advertising executives have recently predicted that display advertising could soon outpace paid search again. This is because the so-called "big money" advertisers (e.g., cars, movies, and packaged goods) have focused their budgets on banner advertising (Baker 2006). Whatever the outcome of the spending race, it is clear that both forms of Internet advertising are important and growing rapidly. Perhaps because paid search is the newer of the two, more research has been conducted into banner ad response and performance metrics than on paid search.

Banner Advertising

Internet users are exposed to banner advertising in the course of browsing websites. One of the interesting features of banner advertising is that it can be purchased in several ways. For example, advertisers can be billed either for *space* on a website, for example, a fixed daily rate, or based on *performance*, that is, impressions served or click-throughs (the user clicks on the ad, taking him to another website). When advertisers buy space to display ads, detailed exposure or click-through measures are not directly connected to campaign costs. The advertiser is billed a fee for the space and time (e.g., hourly or daily) on the site. When located on prime Internet "real estate," such as a banner on a leading portal like Yahoo! or MSN, the costs can be $500,000 a day—about the same as a 30-second spot on a hit series such as CBS's *CSI* (Baker 2006).

Traditional Metrics for Banner Ads

When buying ad space at a fixed rate, a stated goal is often to achieve brand awareness or develop brand associations in consumer memory. As such, the use of survey-based tracking metrics for awareness and brand attitudes is appropriate. In addition to simple tracking metrics, regression-based models can be used to analyze the relationship between changes in banner ad spending and aggregate-level outcomes such as site visitation volume. Ilfeld and Winer (2002) provide a good example of the regression approach to this problem. In their study, they examine the impact of both online and offline advertising on brand awareness, site visitation, and brand equity over time. They found that website visits were significantly related to online advertising spending. They estimated the elasticity of site visits with respect to advertising to be 0.14, a figure within the normal range of advertising elasticities reported in the marketing literature.

While marketing researchers have successfully applied established methods to aggregate-level Internet advertising data, such methods do not take full advantage of the information from the clickstream. In particular, exposure to banner ads, click-through on banner ads, and subsequent purchasing behavior can all be tracked and analyzed at the individual level. This has led to the development of a new set of approaches for assessing performance.

Clickstream-based Metrics for Banner Ads

Because clickstream data enable websites to track the pages viewed by their visitors, it is straightforward to track exposures to the banner ads served on those pages, as well as record whether or not a visitor clicked on a banner ad. Though promising from the viewpoint of having granular-level data, both impressions and click-through are not without problems as metrics. Analogous to the issues in offline advertising, a major concern with exposure counts or number of impressions is whether or not the user in fact processed the ad. Indeed, industry studies show that, on average, only one out of every twelve banner ads is attended to (Baker 2006). In the hopes of addressing this problem, click-through metrics were introduced to discriminate between attended and nonattended ads (Hoffman and Novak 2000; Dahlen 2001).

Chatterjee, Hoffman, and Novak (2003) developed a modeling approach to assess consumer response to banner ad exposure at a sponsored content website. Their approach models the user's banner ad click-through decision using an individual-level binary logit (i.e., a click/no-click formulation). Estimating the model on clickstream data, the authors report a series of interesting findings. For example, there are significant differences across users in click proneness. This suggests that users have intrinsic differences in their interest in banner ads or the extent to which they attend to them. Also, repeated banner exposures do not have the same effect across site visitors; they can increase the click-through rate for less click-prone site visitors but not for others. Additional insights obtained from the model include, for example, the finding that consumers are equally likely to click on banner ads placed early or late in the navigation path through the site.

The click/no-click predictive modeling approach provides a very useful set of individual-level metrics for assessing the effectiveness of banner ads on a website. For example, visitors can be segmented by their propensity to click and by their responsiveness to repeated banner ad exposures. Findings such as these, when applied to a given site, can aid advertisers and site managers in devising banner ad strategies and in potentially targeting ads to different users.

While click-through rates were reported to be as high as 7 percent in 1996, by 2002 the average clickthrough rate had fallen to 0.3 percent (DoubleClick 2003). Drèze and Hussherr (2003) investigated why banner advertising seems to be ineffective based on click-through measures. They found that consumers avoided looking at banner ads. This could imply that pre-attentive processing is taking place, indicating that other metrics (e.g., aided and unaided recall) may be more appropriate for assessing banner advertising.

Instead of looking at click-through, another clickstream-based approach is to focus on purchase outcomes. The idea here is to relate changes in banner ad exposure to changes in individual-level purchase behavior. To do this, Manchanda et al. (2006) propose a hazard modeling approach to capture the relationship between banner advertising and purchase patterns at an e-commerce website. Using detailed clickstream records, they investigate individual purchase timing behavior as a function of banner ad exposure. The authors report that banner advertising can accelerate the purchases made by existing customers (i.e., customers exposed to more banner ads purchased again sooner). They also find that when banner advertising exposures are spread across different websites and pages, the level of response is enhanced, but when the ads use a wider variety of creative executions that response is diminished. These results imply (at least for this site and campaign) that the number of banner ad versions offered in a given campaign be limited but that those ads be spread across a variety of websites and pages.

Like the Chatterjee, Hoffman, and Novak (2003) model for click-through, the Manchanda et al. (2006) modeling approach is useful for gauging individual-level response and targeting possibilities. Manchanda and co-authors also report extensive heterogeneity across users in response to banner ads. While average response to banner ads is low (the mean elasticity of purchase with respect to banner ad exposure is estimated at 0.02), there are nonetheless many users who are quite responsive.

A third approach based on clickstream data analysis is to gauge whether or not exposure to a banner ad alters the future browsing behavior of site visitors. While a banner ad might not trigger a click-through or have a measurable effect on near-term purchase, it is possible that exposure to the ad might affect subsequent consumer search activity on the site. In a working paper, Rutz and Bucklin (2007a) propose a choice model to investigate this question using clickstream data from an automotive website (where extensive consumer search activity takes place). Their approach models the site-browsing choices made by visitors following exposure to specific banner ads for different automotive brands.

Estimating their model, Rutz and Bucklin (2007a) find that banner ads served during a site visit influence the subsequent browsing behavior of the majority of visitors. Three distinct segments of consumers emerged from the analysis: one segment (37 percent of the visitors) responds by seeking more information about the advertised brand, a second segment (23 percent) reacts negatively (moving away from seeking more information about the brand), and a third segment (40 percent) has no significant reaction to the banner ad exposure. The advertising elasticity of browsing choice behavior was found to be approximately 0.2 for the positive reacting segment and –0.05 for the negative reacting segment. The authors also report that ads served in previous site visits had no effect on browsing behavior.

Insights from Consumer Behavior

As noted above, Drèze and Hussherr (2003) investigated why banner ads seem to be relatively ineffective. They used an eye-tracking device to study consumers' attention to online advertising in combination with a survey of Internet users' recall, recognition, and awareness of banner advertising. They suggest that click-through rates are low because many consumers try to avoid attending to banner ads. If so, the authors conclude that processing will be done at the pre-attentive level and measures such as brand awareness and brand recall also should be used.

Cho and Choen (2004) hypothesize that consumers avoid looking at banner advertising on the Internet due to perceived ad clutter, previous negative experience, and perceived goal impediment. They found perceived goal impediment to be the most significant factor underlying banner ad avoidance behavior. The authors recommend that advertisers use customized context-congruent messages to cut down on perceived goal impediment. Following up, Moore, Stammerjohan, and Coulter (2005) investigate the importance of congruity between the website and the ad. They find that congruity leads to favorable effects on attitudes while incongruity leads to better recall and recognition.

Other studies appearing in the consumer behavior literature indicate that established empirical patterns found in the offline world also apply online. For example, Danaher and Mullarkey (2003) report that the duration of ad exposure is an important moderator of banner ad effectiveness. They show that the longer a person is exposed to a web page containing a banner ad, the more likely he or she is to remember that banner ad. The authors also find that user involvement plays a crucial role in understanding banner ad effectiveness. Users in a goal-directed mode are much less likely to recall and recognize banner ads than users who are in a browsing mode. Also analogous to offline advertising, Havlena and Graham (2004) find that banner ad effects decay rapidly over time, at least for the automotive and pharmaceutical categories. This is corroborated by the clickstream findings from Rutz and Bucklin's (2007a) study at an automotive website.

In sum, research into banner advertising has produced a series of metrics, modeling approaches, and empirical insights that should be useful for both academics and practitioners. With respect to clickstream data, regardless of the metric involved (click-through, purchase, or subsequent browsing/search), there is shared consensus that reaction to banner advertising differs significantly across consumers (and may be a negative reaction for many). While average response levels are low, the potential ability to target the more responsive segments means that banner ads are a promising element of the new communications mix. Researchers in consumer behavior have also established why some of this heterogeneity might exist (e.g., goal-directed users versus browsers) and have suggested specific strategies for enhancing the effectiveness of banner ads on the Internet.

Figure 7.1 **Search Results Page Example for Google**

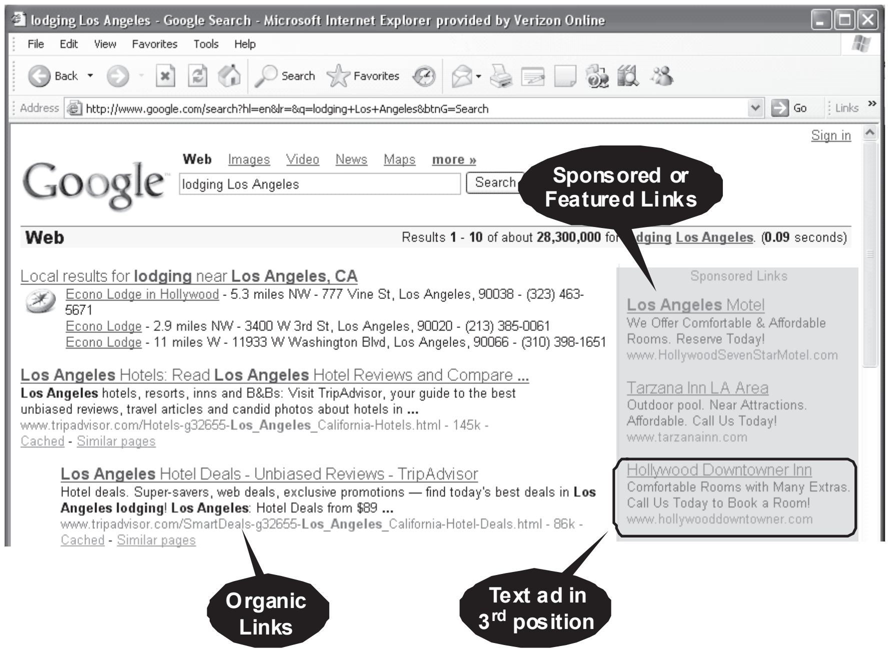

Search Engine Marketing (Paid Search Advertising)

Search engine style information retrieval systems existed long before the Internet; electronic card catalogs, Lexis/Nexis, and other private and corporate search databases have existed for decades. These databases were, for the most part, highly customized expert systems with specific content. Given the content limitations, they were not of widespread interest to the general public. Additionally, access was either highly restricted or expensive. In the last ten years, however, browser-based Internet search engines such as Google and Yahoo! have allowed the general public to readily access information available on the Internet.

This new way of searching for information offers an attractive opportunity to marketers, known broadly as search engine marketing (SEM). While banner ads and traditional offline advertising often intrude on the audience and interrupt activities, search engines offer a less intrusive method as well as an opportunity to offer messages that are congruent with the consumer's goals. This is because search engine traffic originates from voluntary, audience-driven search rather than from forced exposure to a message. Thus, SEM allows marketers to address an audience more likely to be interested in the message and less likely to avoid it. Due to the fairly recent advent of well-organized search engine marketing, most of the academic research on this topic is still emerging and unpublished.

Types of Search Engine Marketing

Marketers use search engines to reach target audiences in two ways. The first is through the natural (also known as "organic") search results returned by the search engine in response to a user query.

The second is through sponsored or paid listings (paid search). Figure 7.1 displays an example of a search results page from Google in which the organic versus sponsored links are identified.

In organic search, the unpaid listings or links are returned and ranked based on the relevancy algorithms used by the search engines. The ranking of a link is based on a combination of factors, including the text content of a web page and external elements such as the number of links to that page at other sites. Clearly, these results are not under the control of advertisers, but they may try to influence them through efforts to improve the ranking of links to the advertisers' website(s).

Efforts to influence organic search results are known as search engine optimization (SEO). SEO seeks to fine tune a website to reflect specific keywords and phrases relevant for the business of the site. The goal of SEO is to try to raise the position of the site's links in the organic listings from a search. This optimization process is technical in its nature (drawing primarily from computer science). To date, we are unaware of any studies in marketing that specifically propose metrics or address the performance outcomes of search engine optimization, and we refer the interested reader to studies available from the industry's professional association for further information (www.sempo.org). It is worth noting that to realize the benefits from SEO, the process needs to be guided by a marketing strategy that identifies target customer segments and the nature of the search activity undertaken on the Internet (e.g., the kinds of terms customers might use in their search queries). Such optimization efforts can also be expensive and may need to be factored into the overall marketing communications budget.

In paid search, an advertiser selects specific search terms (known as keywords) and creates text ads for the search engine to serve (impressions). These ads are displayed either on the top and/or the right side of the search results page and are clearly marked as sponsored or featured listings. To give an example, Figure 7.1 shows the Google search results in response to the query "lodging Los Angeles." Search engines use auction-based methods to determine the position of the sponsored text ads. The selection of the keywords and the bid per keyword can be changed by the advertiser at any time. Advertisers can also limit the amount they are willing to pay for each click by setting a ceiling and a budget.

Many firms now spend a significant part of their marketing budget on paid search. For example, the mutual fund company Vanguard spent $12 million of its $40 million 2005 advertising budget on paid search (Baker 2006). The largest search engine, Google, generated more than $5 billion in net ad revenue in 2006 with their paid search advertising operation called AdWords (Google 2006).

From an advertiser's perspective, paid search ads are not like traditional ads. Advertisers cannot buy their place into the listings for a fixed dollar amount. Unlike with banner ads, for example, companies cannot buy "web real estate." In paid search, advertisers bid what they are willing to pay for a click on a paid search ad. An automated real-time auction algorithm then determines placement and position—for example, first or third—of the ad in the sponsored listings. Thus, paid search is pay-for-performance advertising. Companies do not pay for impressions served, but for actual clicks on their paid search ads. The position of the text ad is determined by the major search engines as a function of the bid price per click and an estimate of click-through. Together, these factors maximize the expected revenue to the search engine. In addition, the relevance of the text ad to the searcher's query is now also being considered in determining the position of sponsored links.

The software provided to advertisers by the search engines allows managers to change the paid search strategy on an almost real-time basis. The selection of the keywords and the bid per keyword can be changed at any time during the day. The software, for example, Google's Adwords, allows advertisers to group keywords, such as Google's Adgroups, and to manage these groups of

Table 7.1

Paid Search Performance, Costs, and Metrics

	Monthly Totals
Impressions	2,281,023
Clicks	14,302
Reservations	518
Cost	$5,107
Conversion rate	3.62%
Cost/click	$0.36
Cost/reservation	$9.86

keywords. Because of the large number of keywords maintained by most advertisers, paid search strategies are often managed based on groupings.

Advertisers can track the results and costs of their paid search campaigns by accessing data from the search engines. This information tells advertisers how many times their text ads have been served in response to a keyword search (impressions), how many times the text ad has been clicked on (clicks), and the cost billed for those clicks. Many companies also subscribe to additional services (usually from third-party vendors) to track whether or not the clicks (which take the searcher to a landing page on the advertiser's site) convert to purchases. From these metrics, the performance of the campaign can begin to be assessed. As an example, the first line in Table 7.1 presents some basic information from a paid search campaign for one firm in the hotel/motel industry in the United States covering a one-month period in 2004.

The table shows the number of impressions, clicks, and lodging reservations accumulated over a month for a managed set of several hundred keywords. It also shows the total cost to the advertiser for the month ($5,107). From this information, several metrics can be readily computed, including the cost per click, cost per reservation, and conversion rate (proportion of clicks that lead directly to a reservation).

An initial assessment of this campaign suggests that it has performed well, generating reservations at a cost per reservation far below the profit from a typical hotel/motel stay at the lodging chain. Such aggregate-level evaluations are a critical part of how paid search advertising is evaluated. However, the aggregation across keywords does not take full advantage of the information available to the advertiser. Specifically, the above information is also available at the level of the individual keyword and/or keyword grouping. If the performance across keywords or across groups differs, then campaign managers have an opportunity to improve returns on spending by altering the keywords selected and/or bidding for those words.

At first glance, one might propose a standard approach to performance evaluation based upon calculating the marginal benefit of spending on each search keyword. Unfortunately, such an approach is infeasible with typical paid search data. Most search keywords have impressions and clicks (costs) associated with them on a daily basis. While a search always leads to an impression, it seldom leads to a click and, even more seldom, a sale. In Table 7.1, the click-through rate is .006 and the conversion rate (reservation given click) is .036, which gives an overall sales (or reservations) per impression rate of .0002. Because the conversion rates in paid search advertising are quite low, many keywords are not associated with any sales whatsoever, even if data are accumulated for several months. This precludes calculating the marginal benefit for a large fraction of keywords. The problem is even worse if managers wish to make decisions about keywords more frequently (i.e., daily or weekly).

To overcome the sparseness problem, managers can look at the performance of groups of keywords and make an assessment at an intermediate level of aggregation. One of the possible categorization schemes for keywords is to group them by generic versus branded. A generic keyword does not contain branded terms (e.g., "Hotels LA"); a branded keyword does (e.g., "Hilton Hotels LA"). These categories are meaningful from competitive, behavioral, and economic performance perspectives. When keywords are divided in this fashion, the advertising cost associated with a sale tied to a generic term is typically much higher than when it is tied to a branded term. However, this does not necessarily mean that the advertiser should favor branded keywords over generic terms. If spillover occurs from generic search to branded search, performance metrics for the keywords need to be adjusted.

In a working paper, Rutz and Bucklin (2007b) investigate the nature and extent of these potential spillover effects in paid search advertising. To do this, the authors model response to paid search advertising using the Nerlove-Arrow Goodwill Model. Exposure to brand-related information served after a search—for example, seeing a paid search text ad or company website—is modeled as advertising that affects advertising stock or "goodwill." The model incorporates the potential effect that this stock has on search-related activity (e.g., impressions, clicks, and sales) over time, thereby revealing whether spillover effects are present.

The authors apply the model to data from a paid search campaign for a major lodging chain (like the data given in Table 7.1). The results show that generic search activity does affect branded impressions, clicks, and reservations. Branded search activity, however, does not affect generic activity, indicating that the spillover effect occurs only in one direction between the two keyword categories. The estimates from the model permit the effective costs for generic and branded words to be adjusted for the spillover effect. This dramatically reduces the costs associated with the generic terms, making them appear far more productive for the advertiser than on the pre-adjusted basis.

Managing paid search campaigns by keyword groups still leaves the detailed word-specific data potentially underutilized. Unfortunately, the sparseness problem noted above means that managers are forced to either drop keywords that do not generate sales or gamble that continuing to buy clicks on those words will pay off in future sales. Alternatively, managers also rely on easy-to-calculate heuristics such as click-through rate (CTR), but there is no guarantee that CTR is predictive of sales conversion.

In a follow-up study to the spillover paper, Rutz and Bucklin (2007c) measure individual keyword performance and address the sparseness problem. The authors develop an approach to estimate conversion rates using a hierarchical Bayes binary choice model estimated across keywords. This approach enables conversion to be based on both word-level covariates and shrinkage across keywords, thereby providing estimates of conversion probabilities even for words that have had no previous sales connected to them. The results show that using the proposed model significantly improves conversion estimates in a holdout sample (versus heuristics or model-free approaches), enabling campaign managers to improve performance.

The metrics and models discussed are geared to aiding managers in campaign evaluation and keyword selection. Paid search advertisers also face several other decisions. For example, they must determine how much to bid for each keyword, how to design the text ads that are served, and how best to design the page that the user lands on when clicking on the text ad. Some research has addressed the bidding decision, but no specific research has yet taken up the design questions. (We speculate that design considerations for text ads and landing pages might share at least some of the issues with traditional print ad design). With respect to the bidding decision, two recent theoretical papers investigate paid search auction mechanisms (Edelman, Ostrovsky, and Schwarz 2007) and paid search advertising as a product differentiation game (Chen and He 2006).

In Edelman, Ostrovsky, and Schwarz (2007), the authors investigate the "generalized second price" auction (GSP). Search engines using the GSP claim that each bidder should bid exactly what they are willing to pay. The authors show that this is a feature of the Vickrey-Clarke-Groves (VCG) auction and not of the GSP auction. They point out that naive buyers who incorrectly assume that they should bid their willingness to pay spend more than necessary while savvy advertisers strategically bid lower.

Chen and He (2006) study a product-differentiation model where consumers are initially uncertain about the desirability of and valuation for different sellers' products. In their model, consumers can learn about a seller's product through a costly search. They find that a seller bids more for placement when his product is more relevant for a given keyword. This results in efficient (sequential) search by consumers and increases total output. Due to their intended theoretical contributions, neither the Chen and He (2006) nor the Edelman, Ostrovsky, and Schwarz (2007) studies examined or modeled actual paid search advertising data.

In sum, much work lies ahead to better understand consumer response to paid search advertising, and to develop and refine appropriate metrics. Managers using the performance and cost data normally provided by the search engines (especially when matched to sales conversion data) can track and assess campaign performance at the aggregate level. To move to the group or keyword level, managers need to be mindful of the potential for spillover among word groups (e.g., generic to branded) and the challenges associated with sparse data on sales conversions for many of the keywords in a paid search campaign.

The Internet as a Social Medium and Online Word-of-Mouth

In this part of the chapter we turn to developments related to the emergence of the Internet as a social medium and the implications for word-of-mouth communications. The Internet has lowered the costs of passing WOM as well as expanded its reach. Blogs, chat rooms, social networks and other forms of user-generated content are growing quickly and user access to them is simple and easy. Search engines also facilitate the rapid collection of thoughts and opinions from people and contexts to which one would otherwise not be exposed. Even commercial websites allow consumers to post feedback, ratings, and comments.

Wikipedia, which, by itself, is a prominent example of this phenomenon, defines social media as a set of online technologies and practices that people use to share content, opinions, insights, experiences, perspectives, and other media (Wikipedia 2007). The online social media (the term also is frequently associated with "Web 2.0") take a variety of forms, from message boards, forums, podcasts, and bookmarks to online communities, wikis, and weblogs. Although the Internet technologies underlying these examples of online social media might be quite distinctive, common to all of them is that they are all primarily user-driven.

In the Internet world, social media marketing are rapidly gaining popularity, and perhaps can be considered one of the forces contributing to e-business's recovery from the dot com crash of the early 2000s. A recent survey conducted by McKinsey (Bughin and Manyika 2007) found that a majority of executives surveyed plan to maintain or increase their investments in technological trends that encourage user collaboration such as peer-to-peer networking, social networks, and web services. Marketing academicians also have long recognized the importance of social media—more and more publications on the topic have begun to appear in the leading academic journals.

In this section we present an overview of some of the recent research developments that may aid in the assessment of these media as marketing communications vehicles. Specifically, we will discuss online communities, aspects of word-of-mouth marketing and, finally, weblogs (or blogs as they are popularly known). As many of these topics are relatively new, a good deal of the research in this area is ongoing or, as yet, unpublished.

Online Communities

In 1995 the first notable social networking (SN) website, Classmates.com, was launched. At the time, no one would have guessed that, twelve years later, the most popular SN site would have tens of millions users and another would have been sold for more than one billion dollars. By most accounts, News Corporation's acquisition of MySpace.com in 2005 marked the beginning of the online SN boom. In 2006, Google paid $1.65 billion for YouTube.com, the online video-sharing community. A series of SN acquisitions are currently taking place in Europe, and several well-established Internet companies such as Microsoft, AOL, and Yahoo! are trying to spin off their own SN sites. SN sites currently attract more than 90 percent of all teenagers and young adults in the United States and have a market of about 80 million members.

The core of an SN site is a collection of user profiles in which registered members can place information that they want to share with others. For the most part, users are involved in two kinds of activities on the SN site: they either create new content by editing their profiles (adding pictures, uploading music, writing blogs and messages, etc.), or they browse through profiles, consuming content created by others (looking at pictures, downloading music, reading blogs and messages, etc.).

The most popular business model for a social networking site is based on advertising. As users browse through a site, a sequence of web pages is served to their screens, each of which may carry ads. SN firms earn revenue from either showing ads to site visitors (impressions) or being paid for each click/action taken by site visitors in response to an ad. SN firms commonly use members' profile information for ad targeting purposes. Accordingly, user involvement with a site (e.g., time spent on the site, number of pages viewed, amount of personal information revealed) can directly translate to firm revenue.

To attract traffic, an SN firm itself cannot do much beyond periodic updates of site features and design elements. The bulk of digital content—the driving force of the site's vitality and attractiveness—is produced by its users. Users, though, are not all created equal. There is a great deal of heterogeneity across community members in terms of the frequency, volume, type, and quality of digital content generated and consumed. A recent article in the *Wall Street Journal* (Holmes 2006) provides anecdotal evidence that a user's interest in a site correlates with content updates produced by contributors. Holmes reports that when a popular blogger left a particular site for a two-week vacation, the site's visitor tally fell and content produced by three invited substitute bloggers could not stem the decline.

The situation with SN sites is more complex: members both produce content and consume content. Furthermore, the average user is likely to be attracted to the SN site by content produced not just by one, but by a number of other users. In contrast to Holmes's story, removing an individual content provider from the site may or may not affect the level of community participation. The community members might choose to maintain their usual activity levels, possibly reallocating their content consumption to other sources on the site. Finally, members also might be attracted to the site because they are motivated to contribute by others' consumption of their own content. In the Holmes example, we might expect a popular blogger to stop producing new posts if people have lost interest in reading them.

From a managerial perspective, understanding who the users are and who keeps the SN site attractive is vital. Who are the key members? What (or who) drives site visitation—and therefore advertising revenues? How should the SN site segment its customers? To manage the site, to make it a better place for users, and to obtain better information for advertisers, the site's management needs to know the roles played by each individual user within the network and whose actions have an impact on whom.

Although still a relatively new area in marketing research, online communities have attracted the attention of many scholars. Dholakia, Bagozzi, and Pearo (2004) studied two key group-level determinants of virtual community participation: group norms and social identity. The authors tested the proposed model using a survey-based study across several virtual communities. They argue that an information-seeker will find the virtual community useful only if he or she can find another participant with the complementary motive of providing that information. To stimulate participation among users, managers need to think in terms of facilitating such exchange by effectively matching participants' complementary motives.

Trusov, Bodapati, and Bucklin (2007) propose a modeling approach to determine the influential members at a social networking website (i.e., who are the members who drive other members' participation?). Using clickstream data from a social networking website, they model members' daily login activities (as a proxy for their activities on the site) and analyze the marginal impact of an individual user's behavior on the other members' activities. The authors find that the top third of members, ranked by influence, account for 66 percent of the total impact on participation. By identifying the key members who are more influential on the network (and are therefore indirectly responsible for more ad revenue), SN firms can selectively target members with retention and participation incentives. In sum, the model provides a method to obtain influence metrics for each member of an Internet social network.

The Trusov, Bodapati, and Bucklin (2007) approach infers influence by examining the actions of one site user versus others over time, as recorded in the clickstream. Another approach to assessing influence comes from the social network analysis (SNA) literature in sociology. In the SNA literature, a social network is defined as a set of actors and a set of relationships (ties) among them (e.g., Scott 1992; Wasserman and Faust 1994). Usually, the importance of an individual actor (in this case, a community member) can be inferred from his or her location in the network (e.g., Iacobucci 1998). In this approach, multiple relational networks can be defined for the same set of actors (e.g., network of e-mail exchanges, network of cross-profile content uploads/downloads, network of cross-profile views and visitations). Since these networks may have different structures, an actor's importance could vary across networks. Van Den Bulte and Wuyts (2007) provide an overview of the applications of SNA to problems in marketing. They discuss, for example, when it is important to consider the network structure in which a consumer resides and how the structure might affect consumer behavior.

Turning from social networks to online forums—another type of online community—researchers have investigated how information is communicated and relationships are formed. For example, Kozinets (2002) developed a new approach to collecting and interpreting data obtained from consumers' discussions in online forums. Labeled by the author as "netnography," the approach can be used by managers to study motivations and consumption patterns of online community members. Using an online coffee community as an example, the author shows how unobtrusive observations of online dialogs may generate some useful marketing research information such as ideas for new products and/or services and opportunities for bundling. Also, frequency of posts and quotations by other members may serve as an indicator for opinion leadership within a community.

Another question addressed by research in online communities is how community members "link" with other individuals within a network. Narayan and Yang (2006) studied a popular online provider of comparison-shopping services, Epinions.com. The authors modeled the formation of relationships of "trust" that consumers develop with other consumers whose online product reviews they consistently find to be valuable. The authors discuss how the formation of "trust" facilitates the flow of information in online communities. Understanding such flows should be able to aid firms in monitoring and, perhaps, facilitating word-of-mouth communications.

Word-of-Mouth Marketing and the Internet

Word-of-mouth marketing has recently attracted a great deal of attention among practitioners. Due to the ease with which consumers can now communicate directly with each other via the Internet, the costs of disseminating WOM have declined and its reach has expanded. In addition, marketers are particularly interested in gaining more understanding of online word-of-mouth as traditional forms of communication may be losing some of their effectiveness (e.g., Forrester 2005). Unfortunately, due to recent development of eWOM, empirical research is remains somewhat limited regarding its effectiveness.

Internet WOM consists of several different types of communication activities. The major categories include viral marketing (creating messages that are designed to be spread widely, as in the spread of an epidemic), referral programs (creating tools for customers to easily refer family and friends), and community marketing (forming or supporting niche communities that are likely to share interests, such as user groups, fan clubs, and discussion forums).[1] In each case, word-of-mouth communications can strongly impact a firm's well-being; however, there has yet to exist a simple set of metrics for assessing these effects, or ones that will guide managers in operating in this domain.

Godes et al. (2005) proposed a useful conceptual framework for thinking about the firm's role in WOM. They argued that the firm can take four possible positions: (1) observer (collect information to learn about its ecosystem), (2) moderator (foster WOM communications), (3) mediator (actively manage social interactions), and (4) participant (play a role in social interactions by actively creating WOM communications). The existing research into eWOM can be classified using these categories, and we discuss several of these below.

The Firm as an Observer of WOM

In an important study bearing on WOM, Godes and Mayzlin (2004) looked at the effect that word-of-mouth communications in online communities had on TV show ratings. The authors proposed that analyzing online conversations could serve as a practical and cost-effective way to measure the effect of WOM. In addition to the volume of WOM communications, the authors also measured the dispersion of WOM across communities. They found that the dispersion measure had strong explanatory power in their model of TV show ratings (i.e., tracking volume alone is not sufficient). These results indicate how tracking WOM on the Internet and modeling its relationship to outcomes can provide an assessment of its impact on product performance. Indeed, the authors recommend that such tracking efforts be instituted early in a product's life cycle so as to capture the critical role of word-of-mouth in the early stages of the life cycle.

Another important venue for word-of-mouth communication on the Internet is product reviews posted by users. The effects of such reviews on book sales were examined in an empirical study by Chevalier and Mayzlin (2006). Looking at book reviews posted on Amazon.com and BarnesandNoble.com within a regression modeling framework, the authors found that an improvement in a book's online reviews led to an increase in sales on that site. They reported that while most reviews tended to be positive and typically resulted in increased sales, a negative review led to a greater impact on sales than a positive one.

Like Godes and Mayzlin (2004) and Chevalier and Mayzlin (2006), recent work by Pai and Siddarth (2007) also empirically links the volume and nature of online WOM to outcome variables. Specifically, Pai and Siddarth use computerized text analysis to quantify the impact that WOM volume and valence (i.e., positive or negative WOM) have on movie box office performance. Their empirical results show that the content of WOM is significantly related to outcomes and that the source of WOM (e.g., movie critics versus consumers) also mediates its impact.

The Firm as a Participant

Using a game-theoretic model, Mayzlin (2006) analyzed the conditions under which a firm should invest in consumer-generated WOM. She showed that word-of-mouth might still be persuasive even if consumers are aware that it could be biased. One of her findings is that low-quality firms spend more resources on promotional chat than do firms with high-quality products.

Godes and Mayzlin (2007) investigated how, if at all, a firm should go about effecting exogenous word-of-mouth. By implementing a large-scale field-based quasi experiment they compared the WOM created by loyal and nonloyal customers and identified the characteristics of influential WOM on sales. One of their key findings is that the marginal impact of WOM created by consumers who have no loyalty to the firm is significantly higher than the WOM created by loyal—and, in some cases, very loyal—consumers. They showed that the most effective (in terms of driving sales) WOM occurs between acquaintances, not friends or relatives, and is created by nonloyal customers. The intuition behind this finding is that members of the nonloyals' social network may be less aware of the product than friends or relatives, and thus represent a more effective group for targeting. The research offers interesting insights on firms' WOM campaign design, such as whom firms should target.

Dellarocas (2005) analyzed how the strategic manipulation of Internet opinion forums affects the payoffs to consumers and firms in markets of vertically differentiated experience goods. He showed that, under certain conditions, forum manipulation can be beneficial to consumers, but in more general cases (e.g., manipulation costs are convex and the number of "honest" consumers' posts is high), such activities end up reducing the profits of all competing firms, as well as overall social welfare.

The Firm as a Mediator

Trusov, Bucklin, and Pauwels (2007) studied the effects of word-of-mouth (WOM) marketing on member growth at an Internet social networking site. Because such sites facilitate and track the electronic invitations sent out by existing members, outbound WOM may be precisely tallied and then linked to the number of new members joining the site (signups). To handle the potential endogeneity among WOM, new signups and traditional marketing activity, the authors used a Vector Autoregression (VAR) modeling approach to quantify the short- and long-term effects of WOM referrals. The authors found that word-of-mouth referrals have substantially longer carryover effects than traditional marketing actions. Specifically, the long-run elasticity of signups to WOM is estimated at close to 0.5, about 2.5 times larger than the average advertising elasticities reported in the literature. The authors also reported that the estimated WOM effect is about 20 times higher than the elasticity for marketing events, and 30 times larger than that of media appearances. Using the estimated Internet display advertising revenue produced by impressions served to a new member, the monetary value of a WOM referral was calculated, yielding an upper bound estimate for the financial incentives a firm might offer to stimulate word-of-mouth.

Blogs

Over the past few years, blogs have become extremely popular sources of news, entertainment, product reviews, and so forth. The blogosphere is a huge, extremely decentralized, and rapidly expanding market. According to Lenhart and Fox (2006) about 12 million American adults keep a blog, while about 57 million read blogs.

In a recent study on the impact of the new product buzz in the consumer packaged goods product category, Nielsen BuzzMetrics (2007) used blog posts as a primary measure of online buzz, among other forms of online WOM activity such as boards and forum postings. While the high correlation between sales and buzz was not surprising, they found that online buzz can actually *cause* sales.

Mayzlin and Yoganarasimhan (2006) analyzed the quality of blogs' signaling mechanism linking them to other blogs under information uncertainty. The naive view of this problem is that one blogger should never link to another (competing) blogger because, by linking, she faces the risk of losing visitors to that competing blog. The authors propose that the nature of the blogosphere actually creates an incentive for bloggers to link, which, in turn, facilitates the consumers' search for information. Using a game theoretical framework they discussed the costs and benefits of linking and showed under what conditions linking to competitors is optimal. In addition, they showed that the number of both incoming and outgoing links can be used as a metric of blog quality.

Conclusion

In the realm of marketing communications, there is perhaps nothing growing faster than the role of the Internet. Spending on Internet advertising is growing much faster than spending on traditional media. And communications by word-of-mouth are accelerating as the Internet enables users to share their reactions and interact with other potential customers to a greater extent—and with greater ease—than ever before. The purpose of this chapter has been to review the recent research developments in marketing that are most relevant to assessing the performance and impact of these communications vehicles.

Internet advertising is composed of two major spending categories: display (or banner) advertising and paid search. To assess display advertising, managers can track exposure data and click-through rates with information derived from the Internet clickstream. Unfortunately, problems exist with both basic metrics because many web users do not attend to the ads on the web pages they browse and click-through rates are extremely small. On the other hand, models have been developed with individual-level data that have been successful in relating click-through rates, purchasing, and subsequent browsing behavior to banner ad exposure. These models offer both the opportunity to calibrate consumer response to banner ads (including an understanding of the segmentation in this response) and to develop targeting strategies. Among the empirical findings on banner ads, a clear consensus has emerged that users are quite diverse in the nature of their response. As such, the low average rates of response to banner ads may mask the fact that some web users are indeed quite responsive to this form of advertising. Part of the reason why some users may shun banner ads is that they are perceived to interfere with goal-directed activity on the Internet.

The second major form of Internet advertising, paid search, may serve as a complement to banner ads in this respect. Because paid search ads are served to users by search engines in response to a specific keyword query, they may be a better fit with consumers who are using the Internet in a goal-directed fashion. While there is less research available on paid search versus banner ads, early work suggests that tracking metrics (currently available to managers based on the performance data typically reported to them) can be refined into response metrics from models that take into account the particular challenges of managing paid search advertising campaigns. In particular, it should be possible for managers to move beyond aggregate-level campaign metrics to assess the performance of keyword groups and/or individual keywords.

In addition to offering new forms of advertising, the Internet also offers users the ability to communicate with greater ease and greater reach. This has drawn renewed attention to the phenomenon of word-of-mouth communication and raised important questions about how firms should deal with it on the Internet. Online communities, social networking sites, online referral programs, and blogs are all examples of the venues in which the Internet facilitates WOM communication. Existing research in the area has shown how electronic records of word-of-mouth (e.g., product reviews, blogs, and so forth) can be connected, via models, to performance outcome variables such as product ratings and sales levels. Emerging work shows promise at identifying the particular members of online communities who are most likely to influence the actions of others, revealing opportunities for firms to manage these interactions and target incentives and creative messages.

While our review of this rapidly evolving subject area is necessarily incomplete and of short shelf life, we hope that the foregoing discussion has been useful in highlighting key findings in the existing and emerging literature and in providing a framework for organizing future thinking and research. There are, of course, other elements of the Internet communications mix available to marketers that we did not have space to discuss. These include, for example, e-mail marketing campaigns and third-party intermediaries (such as shopbots). We hope to include these in our future discussions of this fascinating and fast-growing domain.

Note

1. The interested reader can find a detailed overview of different forms of WOM marketing at the Word-of-Mouth Marketing Association website (www.womma.org).

References

Baker, Stephen. 2006. "The Future of Technology." *Business Week*, March 27.

Bughin, Jacques, and James Manyika. 2007. "How Businesses Are Using Web 2.0: A McKinsey Global Survey." *McKinsey Quarterly Report,* available at www.mckinseyquarterly.com/How_businesses_are_using_Web_20_A_McKinsey_Global_Survey_1913.

Chatterjee, Patrali, Donna L. Hoffman, and Thomas P. Novak. 2003. "Modeling the Clickstream: Implications for Web-based Advertising Efforts." *Marketing Science* 22 (4), 520–41.

Chen, Yongmin, and Chuan He. 2006. "Paid Placement: Advertising and Search on the Internet." Working paper. Leeds School of Business, University of Colorado at Boulder.

Chevalier, Judith, and Dina Mayzlin. 2006. "The Effect of Word-of-Mouth Online: Online Book Reviews." *Journal of Marketing Research* 43 (August), 345–54.

Cho, C. and Cheon, J. (2004. "Why Do People Avoid Advertising on the Internet?" *Journal of Advertising* 33 (Winter), 89–97.

Dahlen, M. (2001. "Banner Advertisements Through a New Lens." *Journal of Advertising Research* (August), 23–30

Danaher, Peter, and Guy Mullarkey. 2003. "Factors Affecting Online Advertising Recall: A Study of Students." *Journal of Advertising Research* (September), 252–67

Dellarocas, Chrysanthos N. 2005. "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms." Working paper. University of Maryland, College Park.

Dholakia, Utpal M., Richard P. Bagozzi, and Lisa Klein Pearo. 2004. "A Social Influence Model of Consumer Participation in Network- and Small-group-based Virtual Communities." *International Journal of Research in Marketing* 21, 241–63.

DoubleClick. 2003. "DoubleClick 2002 Full-Year Ad Serving Trends." www.doubleclick.net, accessed January 2003.

Drèze, Xavier, and F-X. Hussherr. 2003. "Internet Advertising: Is Anybody Watching?" *Journal of Interactive Marketing* 17 (Autumn), 8–23

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars' Worth of Keywords." *American Economic Review* 97 (1), 242–59.

Forrester Research. 2005. "What's The Buzz on Word-Of-Mouth Marketing? Social Computing and Consumer Control Put Momentum into Viral Marketing." www.forrester.com.

Godes, David, and Dina Mayzlin. 2004. "Using Online Conversations to Study Word-of-Mouth Communication." *Marketing Science* 23 (4), 545–60.

———. 2007. "Firm-Created Word-of-Mouth Communication: A Field-based Quasi-Experiment." Working paper. School of Management, Yale University.

Godes, David, Dina Mayzlin, Yubo Chen, Sanjiv Das, Chrysanthos Dellarocas, Bruce Pfeiffer, Barak Libai, Subrata Sen, Mengze Shi, and Peeter Verlegh. 2005. "The Firm's Management of Social Interactions." *Marketing Letters* 16 (3/4), 415–28.

Google. 2006. *Annual Report*. http://investor.google.com.

Havlena, W., and J. Graham. 2004. "Decay Effects in Online Advertising: Quantifying the Impact of Time Since Last Exposure on Branding Effectiveness." *Journal of Advertising Research* (December), 327–32.

Hoffman, Donna, and Thomas Novak. 2000. "Advertising Pricing Models for the Web." In *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*, ed. Brian Kahin and Hal Varian. Cambridge: MIT Press.

Holmes, Elizabeth. 2006. "No Day at the Beach: Bloggers Struggle with What to Do About Vacation." *Wall Street Journal*, August 31, p. B1.

Iacobucci, Dawn. 1998. "Interactive Marketing and the Meganet: Network of Networks." *Journal of Interactive Marketing* (Winter), 5–16.

Ilfeld, Johanna S., and Russell S. Winer. 2002. "Generating Website Traffic." *Journal of Advertising Research* (September–October), 49–61.

Kozinets, Robert V. 2002. "The Field Behind the Screen: Using Netnography for Marketing Research in Online Communities." *Journal of Marketing Research* 39 (1), 61–72.

Lenhart, A. and S. Fox. 2006. "Bloggers: A Portrait of the Internet's New Storytellers." Pew Internet and American Life Project, www.pewinternet.org.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta. 2006. "The Effect of Banner Advertising on Internet Purchasing." *Journal of Marketing Research* 43 (February), 98–108.

Mayzlin, Dina. 2006. "Promotional Chat on the Internet." *Marketing Science* 25 (2), 157–65.

Mayzlin, Dina, and Hema Yoganarasimhan. 2006. "Link to Success: How Blogs Build an Audience by Promoting Rivals." Working paper. School of Management, Yale University.

Moe, Wendy. 2003. "Buying, Searching, or Browsing: Differentiating Between Online Shoppers Using In-Store Navigational Clickstream." *Journal of Consumer Psychology* 13 (1–2), 29–39.

Montgomery, Alan L., Shibo Li, Kannan Srinivasan, and John C. Liechty. 2004. "Modeling Online Browsing and Path Analysis Using Clickstream Data." *Marketing Science* 23 (4), 579–95.

Moore, S., C. Stammerjohan, and R. Coulter. 2005. "Banner Advertiser–Web Site Congruity and Color Effects on Attention and Attitudes." *Journal of Advertising* 34 (Summer), 71–84.

Narayan, Vishal, and Sha Yang. 2006. "Trust Between Consumers in Online Communities: Modeling the Formation of Dyadic Relationships." Working paper, Stern School, New York University.

Nielsen BuzzMetrics. 2007. "The Origin and Impact of CPG New Product Buzz: Emerging Trends and Implications." White paper. June, www.nielsenbuzzmetrics.com.

Pai, Seema, and S. Siddarth. 2007. "The Impact of Word-of-Mouth on Purchase Decisions: The Case of Motion Pictures." Working paper. Marshall School, University of Southern California.

PricewaterhouseCoopers. 2006. "IAB Internet Advertising Revenue Report 2006." May, www.iab.net.

Rutz, Oliver J., and Randolph E. Bucklin. 2007a. "Effects of Banner Ad Exposure on Internet Site Browsing Behavior." Working paper. School of Management, Yale University.

———. 2007b. "From Generic to Branded: A Model of Spillover Dynamics in Paid Search Advertising." Working paper. School of Management, Yale University.

———. 2007c. "A Model of Individual Keyword Performance in Paid Search Advertising." Working paper. School of Management, Yale University.

Scott, John. 1992. *Social Network Analysis*. Newbury Park CA: Sage.

Trusov, Michael, Anand V. Bodapati, and Randolph E. Bucklin. 2007. "Determining Influential Users in Internet Social Networks." Working paper. Robert H. Smith School, University of Maryland.

Trusov, Michael, Randolph E. Bucklin, and Koen Pauwels. 2007. "Estimating the Dynamic Effects of Online Word-Of-Mouth on Member Growth of a Social Network Site." Working paper. Robert H. Smith School, University of Maryland.

Van den Bulte, Christophe, and Stefan Wuyts. 2007. *Social Networks and Marketing*. Cambridge, MA: Marketing Science Institute.

Wasserman, Stanley, and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications.* New York and Cambridge, UK: Cambridge University Press.

Wikipedia. 2007. "Social Media." Available at http://en.wikipedia.org/wiki/Social_media (accessed September 24, 2007).