



Marketing Science Institute Working Paper Series 2012
Report No. 12-100

Social Media Intelligence: Measuring Brand Sentiment from Online Conversations

David A. Schweidel, Wendy W. Moe, and Chris Boudreaux

"Social Media Intelligence: Measuring Brand Sentiment from Online Conversations" © 2012
David A. Schweidel, Wendy W. Moe, and Chris Boudreaux; Report Summary © 2012
Marketing Science Institute

MSI working papers are distributed for the benefit of MSI corporate and academic members and the general public. Reports are not to be reproduced or published, in any form or by any means, electronic or mechanical, without written permission.

Report Summary

Increasingly, businesses are turning to social media as a source of market research. Comments posted on social networking sites, blogs and microblogs, and discussion forums have provided a wealth of data from which marketers have been trying to extract metrics pertaining to the health of their brand. Traditionally, carefully designed surveys have been employed for this purpose. Today, with the proliferation of social media, marketers have turned to social media to “listen in” on the conversations surrounding their brand.

However, key differences exist between traditional research methods and social media listening. A well-designed survey identifies and targets the relevant respondent population; and questions are structured to focus on specific topics of interest, with care taken to avoid any potential response biases. In contrast, social media provides an unstructured and open forum allowing anyone to comment on any topic of interest to them.

In this environment, researchers have identified several factors that influence posted opinions, ranging from venue effects, where the choice of where to post is related to what you post, to social dynamics, where the social interactions in the venue alter the opinions subsequently expressed. As a consequence, metrics based on opinions expressed in social media are often not comparable to those expressed through a well-designed survey.

In this study, David Schweidel, Wendy Moe, and Chris Boudreaux investigate the potential to infer brand sentiment from social media conversations. Their analysis employs data collected from a variety of social media domains. Controlling for various factors that can influence the posted opinion, the authors propose a hierarchical Bayesian regression model and derive a measure of online brand sentiment.

The authors apply their model to data pertaining to a leading enterprise software brand and show how their proposed approach provides an adjusted brand sentiment metric that is correlated with the results of an offline brand tracking survey (correlation = .604). In contrast, a simple average of sentiment across all social media comments is uncorrelated with the same offline tracking survey (correlation = -.002). Their findings show systematic differences in sentiment expressed across different social media venues and across different posters. Additionally, their method provides a tool with which to decompose overall sentiment into an underlying brand sentiment versus sentiment focused on specific products in the brand portfolio or attributes of the brand. The authors further apply their model to a number of brands across different industries and demonstrate potential pitfalls associated with simple average sentiment measures.

Their findings demonstrate the potential for social media to be incorporated into the brand’s research activities; however, these activities must be undertaken with care. Monitoring a single type of venue would not allow managers to distinguish venue-specific factors from the general impressions of the brand. However, firms may be able to infer overall brand sentiment from a broader sample of comments drawn from multiple venues, with consideration given to differences in the comments’ focal attributes and products, posting venue, and customer experiences.

David A. Schweidel is Assistant Professor of Marketing, Wisconsin School of Business, University of Wisconsin, Madison. Wendy W. Moe is Associate Professor of Marketing, Robert H. Smith School of Business, University of Maryland. Chris Boudreaux is Sr. Vice President of Business Integration, Converseon, Inc.

Introduction

With consumers increasingly engaged in online social media, companies have struggled with the question of how to integrate social media into their marketing strategy. Many have treated social media as just another channel with which to promote to and engage with customers. However, the effectiveness of such tactics is largely unknown, in part due to the lack of metrics with which to measure success. Other firms have used social media as a marketing research tool to gauge customer brand sentiment. A recent survey of 200 marketers conducted by Forrester Research found that 88% of respondents monitored online feedback and conversations (Forrester Research 2011). The strong interest in social media monitoring has given rise to a growing industry of “listening platforms” that measure the sentiment expressed through online social media (Hofer-Shall 2010).

This use of social media as a means of obtaining customer insights raises several questions for the marketing research community. In general, we currently have a limited understanding of the behavior related to social media. As a result, researchers have little guidance as to how they should interpret the volumes of comments posted online, leading firms to rely on simplified measures such as the total volume of posted comments or the average sentiment expressed across all posted comments. For example, researchers have monitored the number of tweets as a measure of engagement with products (Rui, Whinston and Winkler 2009) or events, such as the Super Bowl (NYTimes.com 2009) or the 2011 British Royal Wedding (LA Times 2011). Additionally, researchers have tracked aggregate measures of opinion expressed in the text of posted comments to assess viewer reactions to television shows (Kinon 2010) or to predict stock market performance (Bollen, Mao and Zeng 2010).

While these aggregate metrics may offer some guidance for marketers to assess customer engagement, there are several limitations that have prevented marketing researchers from integrating social media listening into their research programs. First, because online environments allow for the posting of open ended and free form content, individuals have the flexibility to comment on anything they wish. As a result, individuals commenting on a given brand may focus their remarks on different attributes of the brand (e.g., customer service vs. reliability of the product itself) or different products in the brand's product portfolio. This stands in contrast to traditionally employed surveys in which researchers elicit responses pertaining to specific topics of interest. With the interest in using social media monitoring to drive market strategy and brand management (Forrester 2011), the implication for social media researchers is that simple metrics based on an aggregation across comments (e.g., average sentiment) can be problematic as they ignore established differences between product- or attribute-specific evaluations and general brand impressions (Dillon et al 2001).

Second, the venue (i.e., website domain) to which an individual posts may be related to the opinion posted. That is, some websites may systematically attract more positive (or more negative) individuals to post depending on a number of factors, such as the format of the venue, the nature of the audience, and various site specific dynamics. As a result, depending on the mix of websites represented in the sample, observed changes in aggregate metrics may simply reflect shifts in the composition of websites in the data sample rather than any underlying shifts in overall perceptions of the brand. Again, offline research methodologies typically control for such sampling biases while online sentiment measures routinely disregard them.

Finally, comments posted online are not always provided by individuals who have the requisite experience needed for an informed evaluation. In traditional offline research, much

care is taken to identify individuals from the population of interest. For example, if a brand were interested in measuring customer satisfaction, it would begin by drawing a sample from the customer base, as these individuals would have the requisite experience. In contrast, methodologies employing social media data tend to examine all posts, regardless of whether the posts are from individuals who have direct experience with a brand's offerings or not. As a consequence, posted opinions may not accurately reflect evaluations of the brand based on customers' experiences. Due to these limitations, brand sentiment metrics constructed from an aggregation of online comments are not necessarily comparable to the measures obtained from traditionally accepted offline methods.

However, the sentiment expressed in online social media can still inform marketing researchers. While there are several factors that influence online posted opinions, these comments are nonetheless affected by the contributor's overall sentiment toward the brand (Buschken, Otter and Allenby 2011). This underlying brand sentiment is distinct from attribute-specific evaluations (Dillon et al 2001), venue effects (Chen and Kirmani 2011), and other venue-specific dynamics (Moe and Schweidel 2011). Therefore, in an effort to derive a metric of underlying brand sentiment from social media conversations, we explicitly model and control for these latter factors to separate their effects from that of the underlying brand sentiment on posted opinions.

Specifically, we consider all posted comments pertaining to a target brand and code the sentiment expressed as negative, neutral or positive. We model posted sentiment as an ordered probit process and separate brand sentiment (a construct similar to the general brand impressions measure proposed by Dillon et al. 2001) from product- and attribute-specific evaluations. Both

of these components of social media sentiment can be of great value to marketers who require customer feedback in the management of their brand image and product portfolio.

Additionally, we consider how the sentiment expressed in a comment varies depending on the website domain to which the comment is posted. We do so by allowing for both domain-specific random effects and systematic differences across venue-types, which we define as the format of posting environment (e.g., blog, micro-blog, discussion forum, ratings and reviews, etc.). After controlling for these effects, we obtain a time-varying measure of brand sentiment that has been adjusted to control for these factors. We further separate this measure into the sentiment of posters who discuss their direct experiences with the brand and those who do not reference a first-hand experience (a proxy to differentiate between customers and non-customers).

To evaluate our approach, we first apply our model to social media data pertaining to a leading enterprise software brand. We compare our proposed brand sentiment measure derived from an analysis of online comments to the results of a traditional offline brand tracking survey that was conducted during the same time period. While we find no relationship between a simple average of the sentiment expressed in posted comments (i.e., ignoring the venue and content of posted comments) and the results of the offline survey (correlation = -0.002), the correlation between our adjusted measure of customers' brand sentiment and the offline survey is $.604$. These results demonstrate that in order for social media monitoring to be used as a marketing research tool, researchers must first explicitly account for factors that can influence expressed online sentiment but may not necessarily reflect overall brand sentiment.

We also apply our modeling approach to three additional social media data sets pertaining to brands in different industries to demonstrate empirical regularities across social

media posting behavior. Across all four social media datasets, our empirical results show that posters referencing a direct experience tend to be more critical than those who do not, a finding consistent with previous research on expertise (Amabile 1983; Schlosser 2005). In three of the four social media datasets considered, sentiment expressed on blogs are more positive compared to other formats, suggesting that solely monitoring blogs may paint an overly optimistic portrait of how the brand is perceived. Our analyses of multiple social media datasets also reveal that the trends in sentiment expressed in individual venue formats fail to track with the overall brand sentiment, highlighting the importance of casting a wide net across venues when monitoring social media.

The remainder of the paper proceeds as follows. We first review research that conceptualizes the individual's posting decision and discuss how various factors influence posting behavior. Next, we describe the social media and survey data we employ from the enterprise software brand. We then detail our analysis of the social media data and our derivation of the proposed brand sentiment measure for a given brand. We discuss the empirical findings and compare the social media based measure of brand sentiment to that obtained from an offline brand tracking survey conducted by the brand in parallel. We further present results from analyses of additional social media datasets from different industries. Finally, we conclude by discussing the implications for social media monitoring as a means of deriving marketing insights.

Online Sentiment

While social media researchers have focused on the measurement of aggregate sentiment online, it is important to understand the individual behavior driving the decision to express an opinion. Therefore, in this section, we review some of the extant research that examines an individual's decision to post an online opinion and discuss factors that influence posted product opinions.

Moe and Schweidel (2011) propose that the posting decision consists of two component decisions, an incidence decision (*whether* to post) and an evaluation decision (*what* to post). Many researchers have focused exclusively on the incidence decision and examined the factors that influence the total volume of online word-of-mouth (Berger and Schwartz 2011; Duan, Gu and Whinston 2008). However, Moe and Schweidel (2011) propose that the incidence and evaluation decisions are inter-related and driven by (1) post-purchase product evaluations and (2) social dynamics in the posting environment. For example, they show that individuals with extremely negative or extremely positive product opinions are more likely to post an opinion online than individuals with moderate opinions, subject to the social dynamics present. Their results integrate the findings of offline research showing that individuals with extremely negative opinions are more likely to engage in word-of-mouth activity (Anderson 1998) with online studies showing a predominance of positive word-of-mouth (Chevalier and Mayzlin 2006).

However, the tendency toward expressing extreme opinions online does not necessarily prevent individuals holding moderate opinions toward the brand from entering the conversation. As brands often represent a portfolio of products (Aaker and Keller 1990), individuals can express an extreme opinion toward a specific product in the brand's portfolio even if they hold moderate opinions toward the brand as a whole. Likewise, brands can be described by a

multitude of attributes (Kirmani and Zeithaml 1993; Zeithaml 1988), and evaluations of specific attributes have been shown to be distinct from general brand impressions (Dillon et al 2001). As such, the opinions provided online may represent only the individual's evaluation on a particular product or attribute and not the underlying sentiment toward the brand.

Several studies have also shown that online opinions can be influenced by audience and venue effects. For example, Schlosser (2005) shows that posters moderate their online opinions in the face of a varied audience, a result consistent with offline studies of multiple audience effects (Fleming et al. 1990). Moe and Schweidel (2011) further demonstrate how, over time, such social dynamics can influence the evolution of opinion in an online environment. As such, since audiences and participants vary across websites, it is likely that the sentiment expressed will vary across social media sites and exhibit differing trends.

Furthermore, a few recent studies have shown that the consumer's choice of where to post is strategic and related to how they evaluate the product being discussed. For example, Chen and Kirmani (2011) show that when an individual's goal is to influence others, that individual will post negative messages in a homogenous venue and positive messages in a heterogeneous forum to more effectively persuade others. Muniz and O'Guinn (2001) show that individuals seek out forums which display beliefs most similar to their own when their objective is to build or strengthen network ties. These studies demonstrate that a poster's evaluation decision is inter-related with his venue-choice decision.

However, despite the documented effects of venue choice on posted opinions, few researchers have controlled for the variation present across venues when constructing sentiment measures, revealing a potential limitation of the extant work on online opinions. Though some researchers have restricted their analysis to a single venue, such as individual newsgroups

(Kozinets 2002), Twitter (Jansen et al. 2009), a retailer's product review environment (e.g., Chevalier and Mayzlin 2006; Moe and Trusov 2011) or a third-party review website (e.g., Duan, Gu and Whinston 2008), this approach does not account for the relationship that may exist between the chosen venue and posted evaluations. Should there exist systematic differences in brand sentiment among venues, any analysis of a single venue would confound venue-specific factors (including venue-specific dynamics) with derived measures of overall brand sentiment.

Finally, posted opinions can differ systematically across posters with their level of experience. Bird, Channon and Ehrenberg (1970) showed that brand perceptions can vary substantially across individuals with different usage frequency, where current customers hold different opinions from those of non-customers. In fact, both online and offline studies have shown that "experts" with greater knowledge and experience with a product are more critical and more likely to express a negative opinion when compared to non-experts (Amabile 1983; Schlosser 2005). Thus, posters who refer to their experience with a product may be more prone to express a negative sentiment compared to posters who do not reference such an experience. Distinguishing between such posters in inferring brand sentiment is of value to firms monitoring social media as a means of acquiring customer feedback (Forrester Research 2011).

The above discussion has outlined several sources of variation to consider when measuring brand sentiment via social media monitoring. Beyond the underlying sentiment toward the brand, extant research has suggested that an individual's posted opinion varies depending on (1) the focal product and/or attribute, (2) the venue to which the comment is contributed, and (3) and whether or not the poster is drawing on his own consumption experience. However, despite the systematic effects documented at the level of the individual poster, many popular social media metrics are based on simple aggregate metrics (e.g., observed

average sentiment) that mask covariate effects and yield a flawed brand sentiment measure. Therefore, in the next sections we discuss an approach to modeling social media data and provide an adjusted brand sentiment metric that tracks a carefully designed and implemented offline survey. We then apply our modeling framework to social media data from other brands and discuss the parallels found across the various brands from different industries.

Data

Our initial analysis involves two datasets from a single enterprise software brand. The first dataset contains consumer comments posted online in a variety of venues. This data was provided by Converseon, a leading online social media listening platform that monitors and records online conversations (Hofer-Shall 2010). Converseon monitors a large sample of website domains and identifies comments pertaining to a target brand. These comments are recorded and the textual content is coded for a random sample of comments. The resulting data set contains approximately 500 postings per month. The online data we use for this analysis spans a 15 month period from June 2009 to August 2010 and contain 7,565 posted comments. These comments were found across over 800 domains that support user generated content.

The textual content of the comments were individually coded by a team of analysts at Converseon. First, comments were coded for sentiment where each comment was identified as positive, negative or neutral. Second, comments were coded to distinguish between those that reference customers' direct experience with the brand and those that were based more on word-of-mouth or other sources of information. For example, comments that provided an anecdote based on personal experience with the brand was identified as a comment from a customer with direct experience. In contrast, comments that referred to a third party source of information (e.g.,

a link to a press release about the brand) were coded as general comments that did not reference a direct experience with the brand. While this is not a perfect measure of the poster's experience with the brand, it does serve as a proxy to differentiate between those who have had direct experience versus those who have not. Finally, the subject of each posted comment is identified. Since the brand represented in this dataset offers a large product portfolio, Converseon's analysts identified the focal product of each comment. Additionally, as the brand can be evaluated along a variety of brand attributes ranging from customer service and support to the technological reliability of its products, the analysts also identified the focal attribute of each comment.¹ In our data, we identify and distinguish between 140 unique products and 59 brand attributes.

From the domains present in our data sample, nine different venue types were identified. Table 1 (Tables follow References throughout) describes each venue format, the number of posted comments they represent in our data, and the proportion of each venue's postings that referenced the poster's direct experience.²

We complement the online data set with a second dataset created from a traditional offline survey conducted by the brand. This survey was administered over the telephone to a sample of 1055 registered customers. The survey was conducted in 10 monthly waves from November 2009 through August 2010, which overlaps with the period during which our online data were collected. The online data, however, contains five additional months of data before the survey started. The survey measured customer satisfaction with the brand using seven separate questions (e.g., "What is your overall opinion about [brand]?" and "How likely would you be to recommend [brand] to a peer or colleague?"). A factor analysis conducted on the seven

¹ Of the 7,565 postings, 2% referenced multiple products and 4% referenced multiple attributes. For these postings, the focal product or attribute was coded as the first mentioned.

² Websites of discussion forums and wikis vary across companies and industries. To provide readers with an example of these domains while maintaining the confidentiality of our data provider, we list the forum and wiki pages for a different software company, Adobe Systems.

individual survey items revealed a single factor with an eigenvalue greater than 1 that explained 65% of the observed variance. Given the single factor on which the survey items load and the high pairwise correlations among the survey items (ranging from .44 to .83), we employ an average response across the seven items to represent our survey-based brand sentiment measure.

In addition to the data from the enterprise software brand, we also have social media data for brands in the financial services, automotive and telecommunications industries.

Unfortunately, offline survey data are not available for these three brands. Therefore, we proceed by first analyzing the social media data for the enterprise software brand and benchmarking the results to the brand's offline survey. We then analyze the social media data for the other three brands to highlight the importance of using a model-based measure of brand sentiment as opposed to observed average sentiment metrics.

Model

Our modeling objective is to measure brand sentiment using the large volume of individual comments in our data sample. Based on the expressed opinions in these comments, we separately identify a latent brand sentiment measure that distinguishes between those comments expressing direct experience with the brand's offerings and those that do not. Our approach provides an adjusted brand sentiment measure that controls the effects of the comment's focal product, focal attribute, and posting venue.

We model the opinion expressed in each comment using a hierarchical Bayesian ordered probit process. For comment i in our dataset, let Y_i denote the posted opinion such that $Y_i = 1$ for negative posts, $Y_i = 2$ for neutral posts and $Y_i = 3$ for positive posts. To estimate the probability associated with the sentiment expressed, we specify $U_i = U_i^* + \varepsilon_i$, where U_i^* is

determined by covariate and random effects, and ε_i is idiosyncratic error. We decompose U_i^* into a venue-specific brand sentiment construct, VS, and comment-specific random effects that allow for variation among comments within a venue for a given month:

$$(1) U_i^* = VS_i + \pi_{p(i)} + \alpha_{a(i)}$$

where $p(i)$ denotes the focal product and $a(i)$ denotes the focal attribute of comment i . We account for heterogeneity across comments related to the focal product and attribute through π and α , respectively, with $\pi_p \sim N(0, \sigma_\pi^2)$ and $\alpha_a \sim N(0, \sigma_\alpha^2)$.

The term VS accounts for variation across the different domains (and consequently different venue formats) to which comments are contributed over time. We define VS to be a function of (1) the general brand impression (GBI) when the comment is posted, (2) time-invariant differences across venues, and (3) temporal variation that occurs within a particular venue format. First, the GBI is specified as a latent construct that varies over time but is common across venue formats. This construct provides the key metric of interest with which to compare the survey-based brand sentiment. Second, differences related to venues are decomposed into a random effect, δ , associated with the website domain of comment i , $d(i)$, and a fixed effect associated with the venue format, $v(i)$. This component of the model allows us to capture both systematic differences across venue formats and unobserved heterogeneity across the large sample of domains to which posters contribute comments. Finally, we allow for variation in the expressed sentiment over time to be specific to the venue format through the term φ .³ This allows for dynamics specific to some venue formats (e.g., social dynamics in ratings and review forums) to influence the venue-specific sentiment measure without necessarily affecting the general brand impression that is common across all venue formats.

³ Due to the sparseness of data and the limited number of observations posted to many of the domains we observe, we assume that this variation is the same across domains of a given venue format.

We define VS separately for those posts that do and do not reference the poster's direct experience as follows:

$$(2) \quad VS_i = \begin{cases} \beta_{v(i)} + \delta_{d(i)} + GBI_{t(i)} + \varphi_{v(i),t(i)}, & \text{Experience}_i = 1 \\ \beta'_{v(i)} + \delta_{d(i)} + GBI'_{t(i)} + \varphi'_{v(i),t(i)}, & \text{Experience}_i = 0 \end{cases}$$

where Experience_i indicates whether or not the i^{th} comment references the poster's experience with the brand's offerings. The coefficient vectors β and β' serve as intercepts for postings referencing a direct experience and those that do not, respectively, that are specific to the venue format. This allows for expressed sentiment to systematically differ across venue formats. We account for heterogeneity across domains of the same venue format through the parameter δ , where $\delta_d \sim N(0, \sigma_\delta^2)$. To account for variation in the general brand impression from month to month that is common across all venues, we assume that $GBI_1 = 0$ and $GBI'_1 = 0$, and we model GBI_t and GBI'_t in subsequent months as:

$$(3) \quad \begin{aligned} GBI_t &= GBI_{t-1} + \theta_t \\ GBI'_t &= GBI'_{t-1} + \theta'_t \end{aligned}$$

for $t = 2, 3, \dots$, where $\theta_t \sim N(0, \sigma_{\theta 1}^2)$ and $\theta'_t \sim N(0, \sigma_{\theta 2}^2)$. That is, GBI and GBI' are each assumed to follow a random walk, allowing for the general brand impression to drift upward or downward relative to the general brand impression from the previous month.⁴

The temporal variation in GBI and GBI' are common across all venue formats. To allow for differences across venue formats in terms of how posted opinions differ over time, we incorporate variables to capture temporal variation that is specific to a given venue format (φ_{vt} and φ_{vt}'), and assume that $\varphi_{vt} \sim N(0, \sigma_{\varphi 1}^2)$ and $\varphi_{vt}' \sim N(0, \sigma_{\varphi 2}^2)$. This specification allows us to differentiate between general changes in underlying brand sentiment (captured by GBI and GBI')

⁴ As an alternative to this model specification, we estimated a series of models in which GBI and GBI' are randomly drawn from normal distributions with mean zero each month and therefore not linked to the previous month. The empirical results under this alternative specification are substantively similar to those found under equation (3).

and more localized fluctuations in opinions that only manifest in certain venue formats. For example, a news event that fundamentally changes underlying brand sentiment will likely result in observable shifts in expressed sentiment across all venues and thus will be captured by GBI and GBI'. On the other hand, a localized event or dynamic specific to a particular venue format has limited implications for the general brand impression and thus will be captured solely through the venue-by-month interactions ϕ_{vt} .

To complete our model specification, we assume that ε_i is drawn from a standard normal distribution. This corresponds to the following ordered probit probabilities:

$$(4) \quad \Pr(Y_i = r) = \begin{cases} \Phi(-U_i^*), & r = 1 \\ \Phi(\mu_{v(i)} - U_i^*) - \Phi(-U_i^*), & r = 2 \\ 1 - \Phi(\mu_{v(i)} - U_i^*), & r = 3 \end{cases}$$

where μ_v are the cutoff criteria (that $\mu_v > 0$) to which the latent sentiment, U^* , is compared. Note that the cutoffs are also specific to the venue format. While the format-specific effects in equations (1) – (3) allow for positive or negative shifts in expressed sentiment across the different venue formats, the format-specific cutoffs allows for the mix of negative, neutral, and positive expressed sentiments to differ. For example, if venue format v is generally less negative, the intercept of VS for venue format v (β_v) will be greater than the intercept for other venue formats while a venue format that encourages more neutral comments relative to positive comments will be captured by a larger venue format -specific cutoff μ_v .

We estimate the model described in equations (1) – (4) using WinBUGS (<http://www.mrc-bsu.cam.ac.uk/bugs/>), which draws from the marginal posterior distributions of the parameters of interest using MCMC. Three independent chains were run for 10,000 iterations. We discarded the first 5,000 iterations of each chain as a burn-in. Convergence was assessed both visually and by Gelman and Rubin's F-test (1992).

Empirical Results

Model fit

We begin our discussion of results by first examining model fit and estimate a series of nested models to evaluate the value of each model component. We begin by estimating a model (Model 1) in which we only consider the month to month variation in the general brand impression, captured through GBI in equation (2), but ignore the remaining sources of variation. That is, in addition to ignoring factors related to the venue format (i.e., $\beta_v = \beta_{v'} = \beta$ for all v), this baseline model assumes that $GBI = GBI'$. We then incorporate random effects associated with specific products, brand attributes and domains (π , α , and δ) in Model 2. Next, in Model 3, we allow for differences in general brand impressions across comments that reference a direct experience and those that do not, relaxing the assumption that $GBI = GBI'$. We further incorporate systematic differences across venue formats (β_v and $\beta_{v'}$) in Model 4. Though the specification in Model 4 allows for differences across venue formats, it assumes that all temporal variation is explained by GBI and GBI' (i.e., $\phi_{vt} = 0$ and $\phi_{vt'} = 0$). We relax this restriction in Model 5 by allowing month to month shifts in sentiment to vary across venues formats.

We compare this set of models by examining DIC and aggregate hit rate in Table 2 (see Tables following References). We calculate the hit rate for each comment as the estimated probability of the observed sentiment, and the aggregate hit rate is the average hit rate across comments.

Comparing our baseline model specification (Model 1), which ignores all differences except for monthly variation, to the full model specification (Model 5), we see a 17.8%

improvement in hit rate from .404 to .476. Coupled with the reduction in DIC, the improvement in hit rate provides support for the inclusion of these model components and highlights the variation in expressed sentiment due to factors not related to general brand impressions.

Based on the model comparison in Table 2, we focus the remainder of our discussion on the results of the fully specified model (Model 5). We start by examining the product, brand attribute and domain effects.

Variation in sentiment across focal products and attributes

An appealing characteristic of the model is its ability to quantify the differences in sentiment across the focal products and brand attributes. This allows the brand manager to isolate how each product in its brand portfolio and how each aspect of their product or service delivery contributes to posted online opinions. In this section, we demonstrate this functionality of the model and examine the posterior estimates for π and α for a selection of products and brand attributes.

For the 20 most popular products in the brand portfolio (based on the volume of comments in our dataset), Figure 1 (Figures follow References throughout) provides the posterior estimates of π . These estimates, in effect, reflect how each product is evaluated relative to the overall brand. For this brand, 18 out of the top 20 products positively contribute to online sentiment. However, products O and T are viewed more negatively relative to the overall sentiment toward the brand. From a brand manager's perspective, these results serve as a red flag and may indicate that some intervention is necessary for these two products.

Figure 2 (Figures follow References) illustrates the variation in sentiment across specific attributes relating to the brand. For illustration purposes, we provide estimates of α for 10

frequently mentioned attributes. While these findings are not generalizable to other brands, we provide these results to demonstrate the ability of our modeling approach to extract sentiment pertaining to specific aspects of the brand's offerings. In this case, seven out of the ten frequently mentioned brand characteristics have a negative effect on expressed sentiment. The exceptions are brand reputation, quality of the product and size of the company. In other words, when posters focus on specific characteristics associated with product performance, their sentiment is more negative. In contrast, when reputation- and trust-related characteristics are evaluated, the sentiment expressed is more positive. For this brand, this indicates that while product functionality may receive critical comments online, the overall brand may be benefiting from a positive halo effect from past successes. Such a result may be cause for concern for the long-term future of the brand if the criticisms of product performance persist and continue to be discussed online.

Differences across venue and in expressed experience

Venue effects on sentiment are captured by both domain specific random effects, δ , and systematic fixed effects associated with various venue formats, β . In much the same way we illustrated product and attribute specific effects, we plot domain-specific effects for 10 frequently occurring domains in our dataset (Figure 3, following references). The results indicate noticeable variation across this subset of domains. For example, the two social network domains represented in the figure have directionally opposite effects on how expressed sentiment deviates from the sentiment expressed in that venue type, where one website domain is attracting more positive opinions than the other. This result underscores the concerns associated with restricting an examination of online sentiment to a single domain.

Differences across domains may result from a number of factors ranging from the design of the site, audience of readers, sponsor of the site, etc. Though these characteristics may be of interest to firms trying to expand their online presence to various social media sites, we treat these differences simply as random effects and focus instead on systematic differences that exist across venue formats. Figure 4 (following references) compares the expressed sentiment across different venue formats for comments that do and do not reference their experience with the brand's offerings.

Across venue types, comments that referenced the posters' direct experiences with the firm's offerings are more negative than those comments that do not reference direct experience. This result is consistent with previous research showing that "experts" tend to be more critical and provide more negative opinions (Amabile 1983; Schlosser 2005). Comparing across venue formats, we see that blogs and social networks are generally more positive than forums and micro-blogs. Due to differences in expressed sentiment across different venue formats, market researchers must be cognizant of the composition of their social media sample when constructing dashboard metrics. Neglecting differences across venue formats may result in the misleading inference that brand sentiment has shifted when the only change may be the proportion of different venue formats represented in the data.

To further illustrate the differences across venue formats, we focus the reader's attention on only the sentiment associated with posts in which the poster's references his/her direct experience. Figure 5 (Figures follow References) plots the monthly means of the posterior distribution of VS for an "average" domain of a given venue format (i.e., $\delta = 0$). We focus on the results from those venues that occur most frequently in our data and account for more than 99% of the observations in our data: blogs, forums, social networks, and micro-blogs. We also

plot in Figure 5 the monthly mean of the posterior distribution for GBI, our adjusted measure of general brand impressions among those expressing direct experience which serves to capture the monthly variation common across all venue formats. We provide the posterior means for the venue-specific sentiment (VS) and GBI in the appendices.

We find that the sentiment expressed in blogs and social networks track closely together through our observation period and is higher than the latent sentiment in micro-blogs and forums. Interestingly, while blogs, social networks, and micro-blogs exhibit an upward trend after month 5 of the observation period, we find a slight downward trend in forums. This is consistent with prior research that show how social dynamics result in a negative trend in posted opinions (Godes and Silva 2011; Moe and Schweidel 2011).⁵

Overall, we see that our estimate of general brand impressions (GBI) generally tracks the venue-specific estimates of sentiment VS for blogs, social networks and micro-blogs. For example, an industry-wide event in which the focal brand was unable to participate due to new guidelines implemented by the organizers (a competitor in the same industry) occurred in month 4. The run-up to this event and the brand's inability to sponsor it may have contributed to the initial decline in sentiment we observe through month 5. After this decline, the adjusted measure of general brand impressions and the venue-specific sentiment measures for blogs, social networks and micro-blogs returned to previous levels and then stabilized. In contrast, the sentiment expressed in forums continues its slow, gradual decline. These differences across venue formats, as well as the departure of the venue-specific sentiment from the common monthly variation captured by GBI, highlight the need to account for venue-related differences in monitoring sentiment online.

⁵ The public data collected from the social network sites more closely resemble micro-blogs among networked individuals. As a result, these comments do not feature the same degree of interactivity present in discussion forums.

To provide a further point of comparison, Table 3 (see Tables following References) provides correlation coefficients between GBI and various observable metrics, such as overall average sentiment, average blog sentiment, average micro-blog sentiment, etc. The first column provides the correlations for the software brand analyzed above. While there is moderate correlation between the observed average sentiment expressed online and our model-based GBI measure, it is clear that simply computing an average sentiment across all comments online (a metric commonly used in industry) is not a substitute for GBI. Furthermore, the correlation between GBI and venue-specific average sentiment measures are highly variable. This highlights the potential risks of focusing on a single venue (or venue format) when measuring brand sentiment. It is also important to note that comments posted to social networks are generally sparse (i.e., in some months, there are no posted comments from which to compute an average social network sentiment). As a result, correlation coefficients are based on fewer months of data as some months have missing data. This highlights an added benefit of the model-based GBI measure which uses Bayesian methods to leverage data observed across months and across venues.

Comparing the Adjusted Measure of Online Sentiment to an Offline Tracking Survey

Having shown the differences between GBI and commonly used social media metrics, we next compare our model-based brand sentiment measures to the software brand's offline tracking survey. Table 4 (see Tables following References) presents the correlations between reported customer satisfaction from the software brand's offline tracking survey and various measures of online sentiment. When we compare the survey results to the average sentiment observed each

month, the correlation is virtually non-existent ($r = -.002$), raising potential red flags for social media researchers tracking online sentiment with aggregate summary statistics.

In contrast, when we compare the offline survey results to our adjusted measure of brand sentiment (*GBI*), the correlation is .604⁶. As this comparison illustrates, differences in sentiment that exist across venues, products or brand attributes must be taken into account when employing social media monitoring. Firms relying on social media monitoring services should be aware of the potential pitfalls of social media metrics that neglect to account for such factors.

We also examine the correlation between the tracking survey and the sentiment expressed in the four venue formats (*VS*) that occur most frequently in our data (blogs, forums, micro-blogs and social networks). While these correlations are higher than the near zero correlation between the survey sentiment and average observed online sentiment, they are lower than the correlation between the survey and *GBI*. This suggests that, while aggregating observed sentiment across multiple venues provides a flawed measure of brand sentiment, measuring sentiment within a single venue does not remedy the issue. Instead, leveraging the information across multiple venues while controlling for factors that systematically influence expressed sentiment provides the best option for an online brand sentiment measure that tracks offline surveys.

Empirical Findings from Additional Social Media Data

To illustrate the generalizability of our framework, we next examine social media data provided by Converseon for businesses in three other industries. We apply the model presented in equations (1) – (4) to data from a U.S. credit card company (25154 social media comments over 12 months), a U.S. auto manufacturer (3751 social media comments over 12 months) and an

⁶ The correlation between the survey-based sentiment and *GBI*' (no direct experience) is .429.

telecommunications provider that operates outside the U.S. (63753 social media comments over 21 months).

Figures 6a, 6b, and 6c (Figures follow References) illustrate the effects of five commonly mentioned brand attributes on the online sentiment expressed for each brand. For the credit card company, many of the attributes relating to its reward programs and sign up bonuses have positive estimates for α above 0, suggesting areas of strength for this credit card provider. In addition, the social media sentiment surrounding fraud and identity theft protection is positive, again pointing to a perceived strength of the firm. In contrast, social media comments focused on fees and interest rates tend to be more negative.

For the auto manufacturer, social media sentiment toward vehicle perceptions and the benefits of reducing the brand line were positive. In addition to these frequently mentioned topics, fuel efficiency (posterior mean of $\alpha = .34$) and R&D (posterior mean of $\alpha = .21$) also had positive estimates for α , suggesting attributes that the firm may want feature in their marketing communications. On the other hand, social media comments focused on management and the overall health of the company, as well as the government bail-out, were more negative.

Lastly, we examine the results for the telecommunications provider. Among those attributes with positive estimates of α are the firm's coverage area and service, as well as the firm's presence in industry news. On the other hand, the comments focused that focused on the brand's reputation, its pricing and the quality of customer service tended to be more negative, suggesting areas in which the firm may consider taking corrective actions. Overall, these analyses illustrate how our modeling approach can be applied to any brand's social media data and how brand managers can utilize the results to diagnose various dimensions of their brand.

Finally, we refer the reader again to Table 3 which provides correlation coefficients between GBI and commonly used average sentiment measures for the three added brands in the last three columns. Similar to the software brand, the average sentiment expressed for the credit card and telecommunication brands are moderately correlated with GBI. However, we also see high variance across venue formats. The observed sentiment for the auto brand, however, is poorly correlated with GBI while observable venue-specific measures are actually negatively correlated with GBI. These results highlight the idiosyncratic nature of observed average sentiment metrics across different contexts. In light of GBI tracking the offline brand survey more closely than the sentiment from individual venue formats in the data from the software brand, these results should give pause to social media researchers who use observed average sentiment measures or focus exclusively on a single venue format.

Discussion

In this research, we conduct a cross-venue analysis of sentiment as inferred from social media comments. In contrast to prior studies on social media that have focused on a single venue or single venue format, our analysis reveals differences in the opinions expressed that exist across venues. Moreover, these are not time-invariant systematic shifts. Rather, the sentiment expressed in different venues shift in distinct ways from month to month. We examine these venue effects for posters that do and do not express experience with the focal brand.

After “backing out” deviations in sentiment that are specific to individual venues, we find that the monthly variation in brand sentiment that is common across venues closely relates to the offline tracking survey administered by the brand. This is in contrast to aggregate measures of observed opinions that are uncorrelated with the survey results, demonstrating the value of the

GBI measure from our modeling framework. Across multiple datasets, we find that the trend in the GBI differs from the sentiment observed in particular venue formats, indicating the importance of monitoring social media across multiple venues. The proposed model also provides an approach for examining specific products in the brand portfolio or specific brand attributes, separately from the global brand. Consequently, social media listening may provide a powerful tool for brand managers interested in assessing specific elements of their brand.

Finally, the current research demonstrates the potential for social media to be incorporated into the brand's research activities. Social media monitoring may offer an economical and timely method from which brand sentiment can be inferred when coupled with automated text analysis (Archak, Ghose and Ipeirotis 2011; Feldman et al 2008; Lee and Bradlow 2011). This would enable firms to supplement their current tracking studies with social media listening. As our research suggests, however, these activities must be undertaken with care. Monitoring a single type of venue would result in the inability to distinguish venue-specific factors from the general impressions of the brand. However, firms may be able to infer overall brand sentiment from a broader sample of comments drawn from multiple venues, provided that the variation across comments due to differences in the comments' focal attributes and products, posting venue and customer experiences are carefully accommodated.

There are a number of directions that remain for future work. While we have accounted for differences in the venues to which social media comments are contributed, we have not investigated the specific characteristics of various social media venues that may influence expressed sentiment. Doing so may provide guidance to brands who are considering incorporating interactive components into their websites. It may also be fruitful to examine how the sentiment observed in different venue formats, as well as the sentiment expressed by

particular individuals, may differentially drive sales or other key performance indicators. While we demonstrate that the adjusted brand sentiment measure is more correlated with the tracking survey compared to the sentiment from any individual venue format, specific venues may have superior predictive power with regards to particular metrics. Finally, while the current research demonstrates the potential for social media monitoring to supplement research programs, further investigation using *both* social media and survey data from a range of categories is essential before market researchers begin to rely exclusively on social media for customer insights.

References

- Aaker, David A. and Kevin Lane Keller (1990), "Consumer Evaluations of Brand Extensions," *Journal of Marketing*, 54 (1), 27-41.
- Amabile, Teresa M. (1983), "Brilliant but Cruel: Perceptions of Negative Evaluators," *Journal of Experimental Social Psychology*, 19 (March), 146-56.
- Anderson, Eugene W. (1998), "Customer Satisfaction and Word of Mouth," *Journal of Service Research*, 1 (1), 5-17.
- Anderson, Eugene W. and Mary W. Sullivan (1993), "The Antecedents and Consequences of Customer Satisfaction for Firms," *Marketing Science*, 12 (2), 125-143.
- Archak, Nikolay, Anindya Ghose and Panagiotis G. Ipeirotis (2011), "Deriving the Pricing Power of Product Features by Mining Consumer Reviews," *Management Science*, 57 (8), 1485-1509.
- Berger, Jonah and Eric Schwartz (2011), "What Gets Talked About? Drivers of Immediate and Ongoing Word-of-Mouth," *Journal of Marketing Research*, forthcoming.
- Bird, M., C. Channon and A.S.C. Ehrenberg (1970), "Brand Image and Brand Usage," *Journal of Marketing Research*, 7 (3), 307-314.
- Bollen, Johan, Huina Mao and Xiao-Jun Zeng (2011), "Twitter Mood Predicts the Stock Market," *Journal of Computational Science*, 2 (1), 1-8.
- Buschken, Joachim, Thomas Otter and Greg Allenby (2011), "Do We Halo or Form? A Bayesian Mixture Model for Customer Satisfaction Data," working paper, Ohio State University.
- Chen, Yu-Jen and Amna Kirmani (2011), "Persuading Others Online: The Consumer as Media Planner," working paper, University of Maryland.
- Chevalier, Judith and Dina Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews," *Journal of Marketing Research*, 43 (3), 345-354.
- Dillon, William R., Thomas J. Madden, Amna Kirmani and Soument Mukherjee (2001), "Understanding What's in a Brand Rating: A Model for Assessing Brand and Attribute Effects," *Journal of Marketing Research*, 38 (4), 415-429.
- Duan, Wenjing, Bin Gu and Andrew B. Whinston (2008), "The Dynamics of Online Word-of-Mouth and Product Sales: An Empirical Investigation of the Movie Industry," *Journal of Retailing*, 84 (2), 233-242
- Feldman, Ronen, Moshe Fresko, Jacob Goldenberg, Oded Netzer, and Lyle Ungar (2008), "Using Text Mining to Analyze User Forums," *Proceedings of the International Conference on*

Service Systems and Service Management (ICSSSM 2008). Melbourne, VIC: IEEE, December 2008.

Fleming, John H., John M. Darley, James L. Hilton, and Brian A. Kojetin (1990), "Multiple Audience Problem: A Strategic Communication Perspective on Social Perception," *Journal of Personality and Social Psychology*, 58 (April), 593–609.

Forrester Research (2011). "Listening and Engaging in The Digital Marketing Age," July, accessed at [<http://www.dellsociallistening.com/>].

Gelman, Andrew and Donald B. Rubin (1992), "Inferences from Iterative Simulation Using Multiple Sequences," *Statistical Science*, 7 (4), 457-472.

Godes, David and Jose Silva (2011), "Sequential and Temporal Dynamics of Online Opinion," *Marketing Science*, forthcoming.

Hofer-Shall, Zach (2010), "The Forrester Wave: Listening Platforms, Q3 2010," Forrester Research.

Jansen, Bernard J., Mimi Zhang, Kate Sobel and Abdur Chowdury (2009), "Twitter Power: Tweets as Electronic Word of Mouth," *Journal of the American Society for Information Science and Technology*, 60 (11), 2169-2188.

Kinon, Cristina (2010), "Creators of 'Glee,' 'Modern Family,' track fan reaction via Twitter," NYDailynew.com, February 3, accessed at http://articles.nydailynews.com/2010-02-03/entertainment/27055284_1_creator-and-executive-producer-pregnancy-twitter.

Kirmani, Amna and Valarie Zeithaml (1993), "Advertising, Perceived Quality and Brand Image, in David A. Aaker and Alexander Biel, eds., *Brand Equity and Advertising*, Erlbaum: NJ, 143-162.

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.

Kuksov, Dmitri and Ying Xie (2010), "Pricing, Frills, and Customer Ratings," *Marketing Science*, 29 (5), 925-943.

LA Times (2011), "Royal Wedding: Twitter, ABC News team up to track chatter online," April 28, accessed at [<http://latimesblogs.latimes.com/technology/2011/04/royal-wedding-abc-news-twitter-look-to-curate-chatter-on-prince-william-and-kate-middletons-big-day.html>].

Lee, Thomas Y. and Bradlow, Eric T. (2011), "Automatic Construction of Conjoint Attributes and Levels From Online Customer Reviews," *Journal of Marketing Research*, 48 (5), 881-894.

Moe, Wendy W. and David A. Schweidel (2011), "Online Product Opinions: Incidence, Evaluation and Evolution," *Marketing Science*, forthcoming.

Moe, Wendy W. and Michael Trusov (2011), "The Value of Social Dynamics in Online Product Ratings Forums," *Journal of Marketing Research*, 48 (3), 444-456.

Muniz, Albert M., Jr. and Thomas C. O'Guinn (2001), "Brand Community," *Journal of Consumer Research*, 27 (4), 412-432.

NYTimes.com (2009), "Twitter Chatter During the Super Bowl," accessed at [http://www.nytimes.com/interactive/2009/02/02/sports/20090202_superbowl_twitter.html].

Rui, Huaxia, Andrew Whinston and Elizabeth Winkler (2009), "Follow the Tweets," *Sloan Management Review*, November 30, accessed at [<http://sloanreview.mit.edu/executive-adviser/2009-5/5152/follow-the-tweets/>].

Schlosser, Ann (2005), "Posting Versus Lurking: Communicating in a Multiple Audience Context," *Journal of Consumer Research*, Vol. 32, September, pp. 260-265.

Zeithaml, Valarie A. (1988), "Consumer Perceptions of Price, Quality, and Value: A Means-End Model and Synthesis of Evidence," *Journal of Marketing*, 52 (3), 2-22.

APPENDIX: Parameter Estimates

Table A1. Mean (SD) across draws for GBI and VS for posters expressing direct experience (assuming $\delta = 0$)

Month	GBI	Blog	Forum	Micro-Blog	Social Network
1	0	2.49 (0.17)	1.14 (0.10)	1.02 (0.24)	2.33 (0.44)
2	-.12 (0.16)	2.26 (0.17)	.82 (0.10)	1.15 (0.21)	2.05 (0.48)
3	-.30 (0.22)	1.53 (0.22)	.91 (0.12)	1.12 (0.19)	1.87 (0.48)
4	-.31 (0.20)	1.95 (0.22)	.95 (0.13)	.84 (0.20)	1.86 (0.49)
5	-.46 (0.23)	1.27 (0.17)	.85 (0.11)	.72 (0.22)	1.72 (0.49)
6	-.36 (0.20)	1.71 (0.19)	.88 (0.11)	.87 (0.22)	1.81 (0.48)
7	-.24 (0.21)	1.86 (0.19)	1.04 (0.11)	.98 (0.21)	1.99 (0.47)
8	-.18 (0.19)	2.08 (0.18)	.62 (0.13)	1.46 (0.22)	2.03 (0.43)
9	-.21 (0.20)	1.76 (0.21)	.95 (0.12)	1.30 (0.24)	1.94 (0.43)
10	-.26 (0.21)	1.65 (0.18)	.64 (0.12)	1.38 (0.25)	2.01 (0.40)
11	-.30 (0.23)	1.74 (0.29)	.80 (0.10)	.94 (0.31)	1.87 (0.49)
12	-.13 (0.22)	2.11 (0.29)	.69 (0.11)	1.71 (0.29)	2.06 (0.49)
13	-.19 (0.22)	1.80 (0.30)	.70 (0.10)	1.32 (0.25)	1.98 (0.49)
14	-.20 (0.21)	2.08 (0.27)	.69 (0.12)	1.35 (0.23)	1.90 (0.42)
15	-.26 (0.22)	1.76 (0.24)	.76 (0.19)	1.40 (0.21)	1.65 (0.39)

Table A2. Mean (SD) across draws for GBI' and VS for posters not expressing direct experience (assuming $\delta = 0$)

Month	GBI'	Blog	Forum	Micro-Blog	Social Network
1	0	2.09 (0.14)	1.34 (0.18)	2.02 (0.21)	2.21 (0.44)
2	-.17 (0.16)	2.18 (0.13)	1.18 (0.21)	1.44 (0.19)	2.04 (0.43)
3	-.30 (0.18)	2.04 (0.13)	1.02 (0.21)	1.27 (0.19)	1.92 (0.44)
4	-.28 (0.19)	2.03 (0.14)	1.14 (0.23)	1.50 (0.20)	1.83 (0.42)
5	-.62 (0.19)	1.52 (0.10)	.73 (0.23)	1.16 (0.19)	1.60 (0.44)
6	-.59 (0.19)	1.51 (0.12)	.72 (0.22)	1.15 (0.20)	1.66 (0.44)
7	-.57 (0.19)	1.65 (0.12)	.70 (0.24)	1.10 (0.20)	1.64 (0.40)
8	-.32 (0.19)	2.10 (0.14)	1.00 (0.22)	1.37 (0.22)	1.95 (0.40)
9	-.36 (0.19)	2.16 (0.13)	.80 (0.22)	1.35 (0.21)	1.73 (0.39)
10	-.32 (0.20)	1.92 (0.13)	.93 (0.24)	1.25 (0.25)	1.84 (0.39)
11	.19 (0.21)	2.66 (0.25)	1.50 (0.23)	1.92 (0.24)	2.49 (0.44)
12	.16 (0.20)	2.53 (0.27)	1.29 (0.23)	2.05 (0.20)	2.34 (0.44)
13	.15 (0.20)	2.45 (0.25)	1.32 (0.22)	1.86 (0.21)	2.47 (0.43)
14	.13 (0.20)	2.54 (0.19)	1.39 (0.26)	1.72 (0.22)	2.44 (0.42)
15	.04 (0.21)	2.42 (0.16)	1.23 (0.27)	1.77 (0.22)	2.19 (0.40)

Venue Format	Illustrative Website	Frequency	Direct Experience
Discussion Forum	forums.adobe.com	2728	93%
Micro-blog	twitter.com	2333	37%
Blog	wordpress.com	2274	23%
Social Network	linkedin.com	155	40%
Mainstream News	cnbc.com	36	3%
Social News	digg.com	19	47%
Wiki	adobe.wikia.com	10	50%
Video	vimeo.com	6	0%
Review sites	epinions.com	4	25%

Table 1. Frequency of Venue Formats for Software Brand

Model	Description	DIC	Hit rate (Improvement over Model 1)
1	Monthly differences	14934	.404
2	+ random effects from product, attribute, and domain	13946	.455 (12.6%)
3	+ experience difference	13773	.460 (13.9%)
4	+ venue main effect	13549	.470 (16.3%)
5	+ venue x time interactions	13454	.476 (17.8%)

Table 2. Model Performance for Software Brand

Measure	Software	Credit Card	Auto	Telecom
Observed Aggregate Sentiment	0.532	0.632	0.285	0.681
Observed Blog Sentiment	0.480	0.630	-0.058	0.903
Observed Forum Sentiment	-0.143	0.525	-0.187	0.561
Observed Micro-Blog Sentiment	0.691	0.481	-0.725	0.808
Observed Social Network Sentiment	0.721	-0.376	-0.012	0.112

Table 3. Correlation between *GBI* and Observable Social Media Metrics

Measure	Correlation with Offline Survey
Observed average sentiment	-.002
VS – blogs	.451
VS – forums	-.339
VS – micro-blogs	.451
VS – social networks	.218
GBI	.604

Table 4. Correlations between Social Media Metrics and Offline Survey for Software Brand




Figure 1. Effect of Focal Products on Brand Sentiment for Software Brand




Figure 2. Effect of Focal Attribute on Brand Sentiment for Software Brand




Figure 3. Differences in Brand Sentiment Across Domains for Software Brand

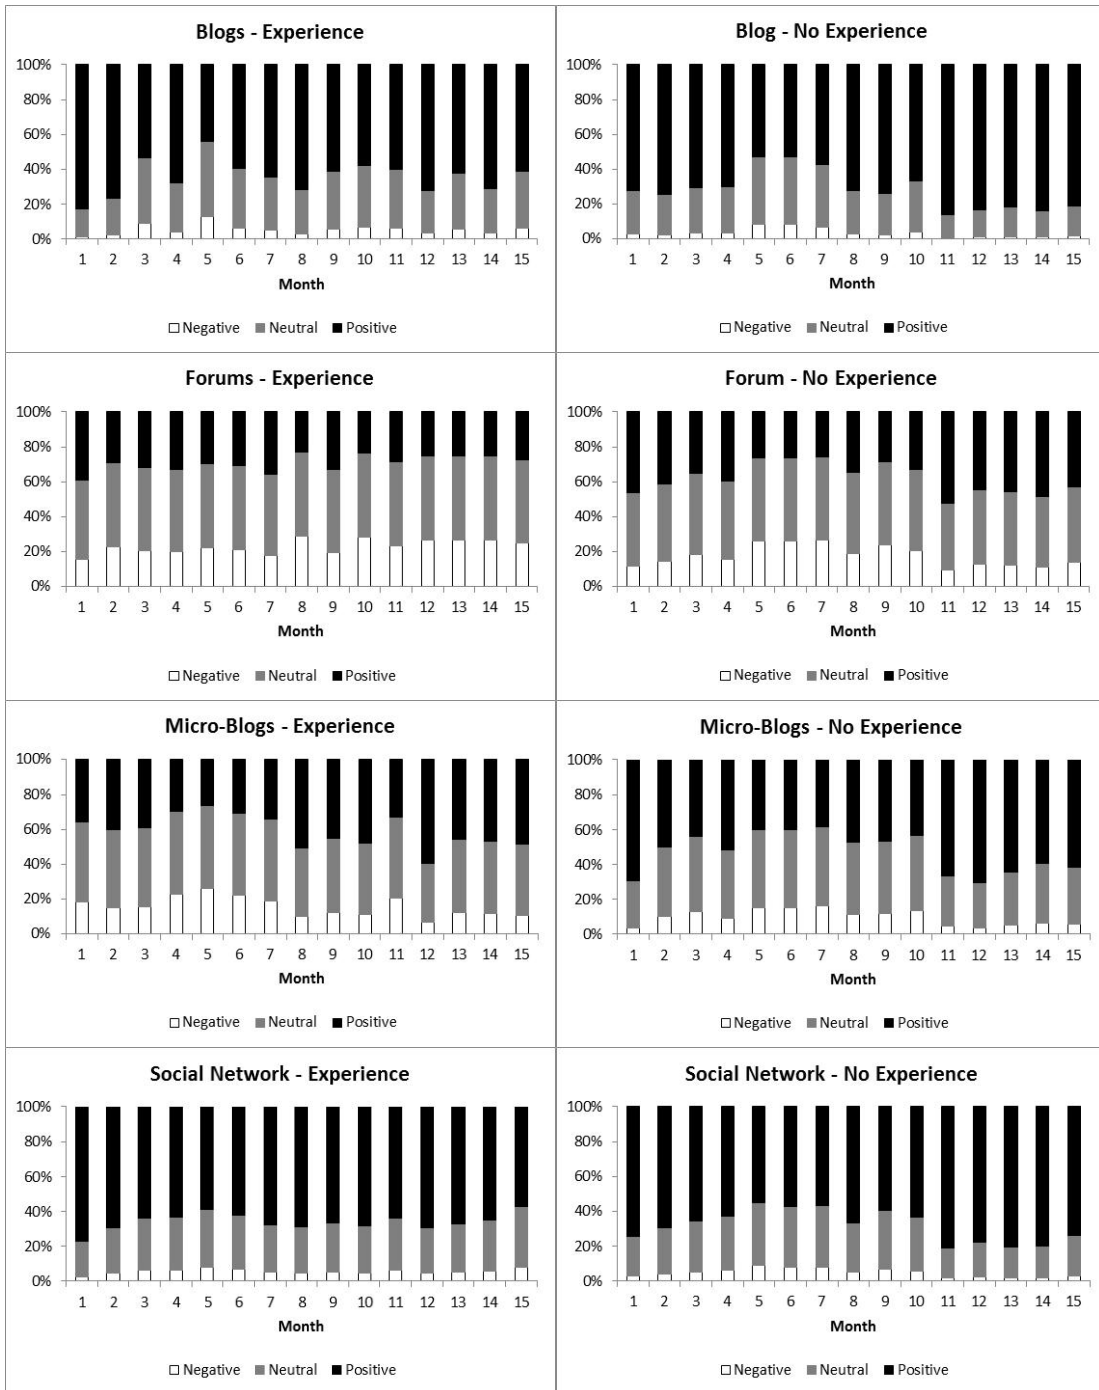


Figure 4: Expected Brand Sentiment Distribution across Venues for Software Brand

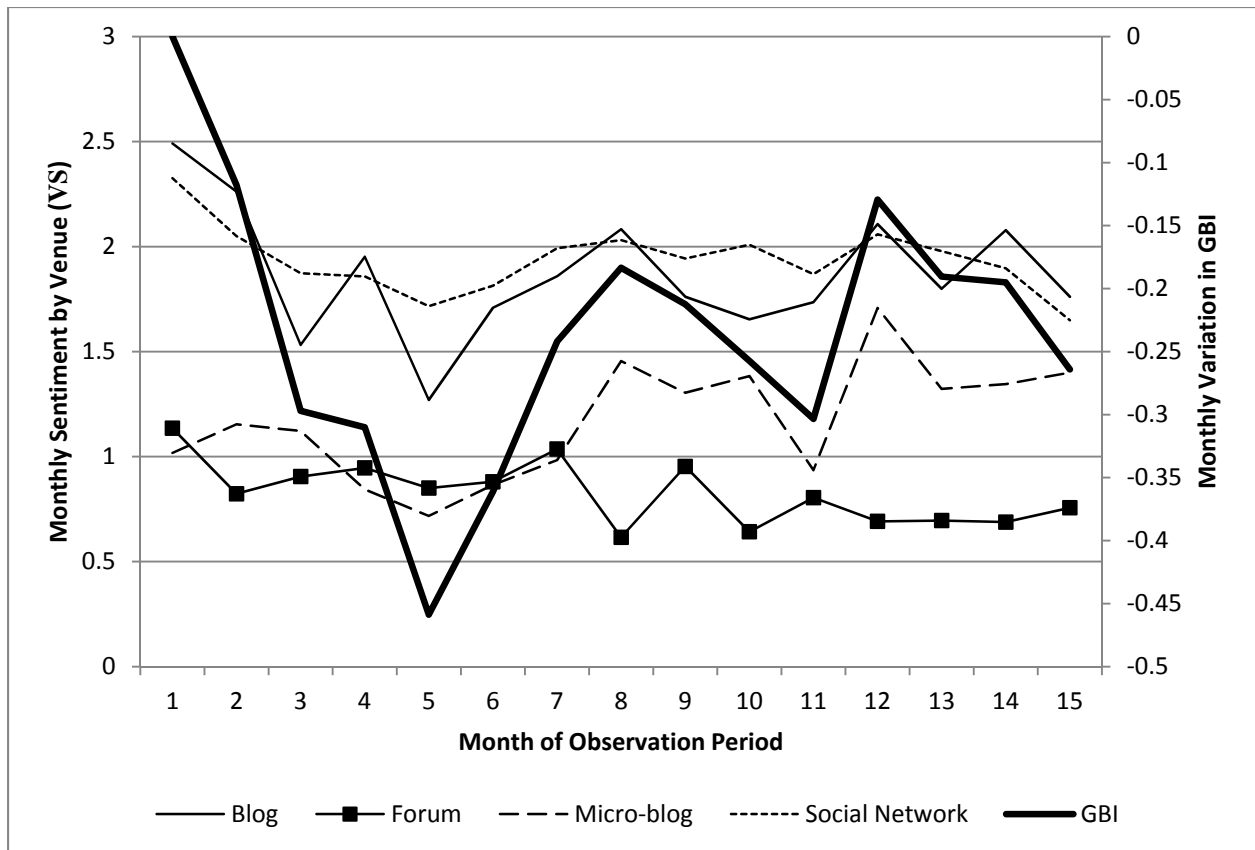


Figure 5. Sentiment Over Time for Software Brand




Figure 6a. Effect of Focal Attribute on Brand Sentiment for Credit Card Brand




Figure 6b. Effect of Focal Attribute on Brand Sentiment for Automobile Brand




Figure 6c. Effect of Focal Attribute on Brand Sentiment for Telecommunications Brand