On the Optimal Marketing Aggressiveness Level of C2C Sellers in Social Media: Evidence from China*

Xu Wang[1], Bart Baesens[2] and Zhen Zhu[†1,3]

[1]IMT School for Advanced Studies Lucca, Italy
[2]Faculty of Economics and Business, KU Leuven, Belgium
[3]Centre for Business Network Analysis, University of Greenwich, United Kingdom

Abstract

Social media has become a widely used marketing tool for reaching potential customers. Because of its low cost, social media marketing is especially appealing to customer-to-customer (C2C) sellers. Customers can also benefit from social media marketing by learning about products and by interacting with sellers in real time. However, a seller's marketing microblogs may backfire on her for dominating the social space. Defining the marketing popularity as the average number of likes each seller receives per marketing-related microblog and defining the marketing aggressiveness level as the proportion of her marketing-related microblogs, this paper empirically quantifies the optimal level of marketing aggressiveness in social media to achieve the maximum popularity. We gather the data from China's largest microblogging platform, Sina Weibo, and the sellers in our sample are from China's largest C2C online shopping platform, Taobao. We find that the empirical relationship between the marketing aggressiveness level and the marketing popularity follows an inverted U-shape curve, where the optimal level is around 30%. In addition, we find a saturation effect of the number of followers on marketing popularity after it reaches around 100,000. Our finding implies that social media marketing should not overlook customers' social needs. Our measure of marketing

*We thank the helpful comments from the participants of the seminars at IMT School for Advanced Studies Lucca, KU Leuven, and SCERI of Stockholm School of Economics, and the insightful discussions with Michelangelo Puliga and Massimo Riccaboni. All the remaining errors are our own. The work of Z.Z. was supported by the "ViWaN: The Global Virtual Water Network" project, the National Research Project (PNR) CRISIS Lab, and the Lars-Erik Thunholms Foundation for Scientific Research.

†Email: zhen.zhu@imtlucca.it, corresponding author.

aggressiveness provides a dynamic business metric for practioners to monitor so as to improve their marketing and managerial decision making process.

Keywords: *Marketing, Aggressiveness, Social Media, Electronic Commerce, Microblog, Popularity*

1 Introduction

Built on the Web 2.0 technology, social media is the kind of applications that allow the creation and exchange of User Generated Content (UGC) on the Internet. It includes collaborative projects such as Wikipedia, social networking sites such as Facebook, microblogging platforms such as Twitter, and visual content communities such as YouTube. More recently, Mobile Web 2.0 (i.e., Web 2.0 evolution with mobile devices) expands the scope of social media to an unprecedented scale (Kaplan & Haenlein 2010). For example, an increasing number of merchants have employed social media as a marketing tool in electronic commerce. They can promote their goods by simply posting a message, usually containing pictures, links, and a short description of them. Compared with traditional marketing channels, social media has provided a more efficient and economic way for sellers to reach potential customers. Moreover, sellers can get feedback in real time and gain a better understanding of customers' demands through their likes and comments, i.e., the popularity of microblogs (De Vries et al. 2012).

Not only sellers benefit from social media marketing, but customers gain gratifications as well. By reading the microblogs, they can get information of the products

they need. However, customers use social networks also for social connection and entertainment (Lin & Lu 2011). While they appreciate a moderate level of marketing communication on products, customers may get annoyed when a marketer is too aggressive in promoting her products.

Previous literature has mainly focused on how consumer perception is affected by social media marketing (Mangold & Faulds 2009, Kaplan & Haenlein 2010). Some studies have demonstrated people's negative attitude towards aggressive social media (and online) marketing by making survey (Grant 2005, Akar & Topçu 2011) and investigated the factors driving customers' engagement with marketing information (Bauer et al. 2005, Chu & Kim 2011, Chu & Choi 2011, Tsai & Men 2013). However, the level of marketing aggressiveness has not been explicitly measured in the above studies.

There is also a paucity of literature on C2C sellers' behaviour in social media. Prior studies on social media marketing are by and large confined to business-to-customer (B2C) (Kumar & Mirchandani 2012, Rapp et al. 2013, Taylor et al. 2011), business-to-business (B2B) (Michaelidou et al. 2011, Swani et al. 2014, Wiersema 2013), and electronic word-of-mouth (eWOM) between customers (Chu & Kim 2011, Chu & Choi 2011, King et al. 2014). This is mainly due to the practical difficulty of studying C2C business in social media. On the one hand, unlike firms that can be searched by their brand names, C2C sellers have little additional information to be identified in social media. On the other hand, it is rare that a large number of C2C sellers from the same e-commerce platform use the same social media platform to promote their products, which makes the data collection a formidable, if not impossible, task.

This paper attempts to fill the gap by analysing the data of Taobao (China's largest C2C e-commerce platform, similar to eBay and Amazon) sellers on Sina Weibo[1] (China's largest microblogging platform, similar to Twitter[2]). We overcome the difficulty of studying C2C business in social media thanks to the collaboration between Taobao and Sina Weibo. On August 5, 2013, Sina Weibo, together with Taobao, released a new module specifically designed for Taobao sellers. The new module gives the verified Taobao sellers additional capabilities (compared with other regular Sina Weibo users) to promote their merchandise. More importantly, it grants the verified Taobao sellers an identity of "Tao," which is highlighted in their Sina Weibo profiles.[3] Moreover, in January 2014, Sina Weibo cooperated with Alipay (Taobao and Alipay are both subsidiaries of Alibaba Group) to launch a new platform called Weibo Payment, making money transfer much easier for Sina Weibo users. These policies have encouraged more Taobao sellers to create accounts on Sina Weibo and to make full use of this marketing channel. The gathering of Taobao sellers on Sina Weibo gives us a unique opportunity to analyse their marketing behaviour and to provide marketing strategies for them. Such marketing strategies would be of great importance for C2C sellers since they cannot afford to market through the traditional channels such as TV, newspapers and magazines.

We first identify 52,187 Taobao sellers on Sina Weibo and collect their microblogs

[1]Some recent studies on Sina Weibo include Guan et al. (2014) and He & Song (2015).

[2]Sina Weibo implements most features of Twitter. Some differences include the use of hashtag (Sina Weibo uses #HashName#) and the charater limit (from 2016, Sina Weibo ended the 140-character limit) that are mostly due to linguistic and cultural reasons.

[3]The minimum requirement for applying for an identity of "Tao" is that the virtual store owner in Taobao should have the level of credibility of at least "one diamond," see `http://help.weibo.com/newtopic/taobao/list/1770/1772` for more details.

in November 2014. For the 12,744 sellers who add the links of their Taobao shops, we further track their microblogs from July to October in 2016. For each seller, we use the proportion of her marketing-related microblogs to measure her marketing aggressiveness. To define whether a microblog is about marketing, we employ different machine learning algorithms to a training set of 5,000 microblogs which we manually labelled. The Multinomial Naive Bayes classifier has the best performance and is thus used to classify all the rest microblogs. Defining the marketing popularity as the average number of likes a seller receives per marketing-related microblog, we conjecture that the relationship between the marketing aggressiveness level and the marketing popularity follows an inverted U-shaped curve. By performing different models to explore their relationship, we find that the linear regression model using Yeo-Johnson transformation of the number of followers has the best performance. After multiple tests, we empirically confirm that there is an inverted U-shape relationship between the marketing aggressiveness level and the marketing popularity. Specifically, the optimal proportion of marketing microblogs is around 0.3.

Our finding highlights the tension between customers' social and commercial needs. A good practice of social media marketing should strike a balance between the two. Acknowledging the dynamic and networked nature of social media, a burgeoning literature aims at constructing a set of social media metrics for marketing and management practioners (Peters et al. 2013). Our measure of marketing aggressiveness provides a strong candidate for such a set of social media metrics. Being informed about the aggressiveness level in real time, the sellers can improve their decision making process and adjust their marketing input as the aggressiveness level becomes too

high or too low. Moreover, our empirically justified optimal level of aggressiveness is a concrete number, or at least a limited range, for social media marketer to keep in mind. However, the optimal level should not be interpreted as a "silver bullet" or be targeted mechanically. We rather suggest that customers' social and commercial needs do not necessarily exclude eath other and new forms of marketing content can be developed to relieve the tension. For instance, the marketing microblogs can be creatively combined with social and personalised elements which customers can modify, share, or simply consume with a better experience.

The remainder of the paper is structured as follows. In Section 2 we present the conceptual framework and hypothesis. Section 3 discusses the design of the study, the algorithms to classify the microblogs and the models to explore the relationship between the marketing aggressiveness level and the marketing popularity. Section 4 provides the estimation results and the analysis before Section 5 concludes the paper.

2 Conceptual Framework

Our focus in this paper is the relationship between the marketing popularity and the marketing aggressivenss level. Below we first provide the clear definitions for the relevant variables before stating the hypothesis of the relationship.

2.1 Marketing Popularity

Social media has effectively cultivated UGC (Gangadharbatla 2008). New ideas to monetize social networks and UGC have encouraged more and more sellers to create public accounts on social media and to increase the number of followers (Ballings et al. 2016). By leveraging the power of eWOM, sellers can disseminate the news of their products to a large number of potential customers. Moreover, customers can interact with sellers and publicly state their opinions by liking, commenting, and forwarding (retweeting) their microblogs.

Previous studies use the number of likes, comments, and forwards to measure the popularity of microblogs (De Vries et al. 2012, Yu et al. 2011). Since a user can comment or forward a microblog as many times as she wants while she can only like a microblog once on Sina Weibo, we adopt the number of likes as the measure of popularity in our study. More importantly, we consider the likes that sellers receive only for the microblogs related to marketing behaviour because they can directly reflect the purchase intention of potential customers.

Therefore, our measure of marketing popularity is the popularity in the social media space and should not be confused with other marketing outcomes such as sales performance. In fact, Schivinski & Dabrowski (2014) argue that high level of popularity in social media has a positive influence on brand equity and brand attitude, which in turn shows a positive influence on purchase intention. Hence, the popularity of microblogs is of crucial importance to sellers. They must have a good knowledge of the factors affecting the popularity of their microblogs and devise strategies accord-

ingly.

2.2 Marketing Aggressiveness Level

Our paper is also related to the longstanding literature of uses and gratifications (U&G) theory (Katz et al. 1973, Eighmey & McCord 1998, Ruggiero 2000). In principle, potential customers can gain gratifications by interacting with both marketing microblogs (i.e., being informed or educated about products and services) and non-marketing microblogs (i.e., being connected in social life). We define the marketing aggressiveness level as the proportion of marketing microblogs. Given a moderate aggressive level of marketing, potential customers can gain both social and business benefits and are more likely to interact with sellers. However, if the microblogs are solely focused on business, followers may get annoyed because they use social networks also for social connection and entertainment (Lin & Lu 2011). On the other hand, the microblogs cannot be solely focused on regular social interactions due to the obvious marketing motivation. Therefore, either extreme of the spectrum will likely decrease customers' gratifications and discourage them from liking the marketing microblogs and the optimal marketing aggressiveness level to attain the maximum marketing popularity should be somewhere in the middle. We finally state our hypothesis as below.

Hypothesis: The relationship between the marketing aggressiveness level and the marketing popularity, as we have defined and measured above, follows an inverted U-shape curve.

It follows from the hypothesis that there exists an optimal marketing aggressiveness level to achieve the maximum popularity. We test the hypothesis by regressing the marketing popularity on the marketing aggressiveness level (and its squared term) along with other control variables. The hypothesis is supported if a significant inverted U-shape curve exists between the two variables in a statistical sense.

2.3 Control Variables

We control other variables that may affect the marketing popularity such as whether the seller's identity has been verified, the seller's gender, her number of followers and her average number of pictures posted per marketing microblog.

The verification of identities provides a signal of trust and reputation. The importance of trust in e-commerce has long been emphasized in previous studies (Gefen 2000, Gefen & Straub 2004, Hoffman et al. 1999, Teo & Liu 2007). Trust is a major factor that affects the prosperity and success of e-commerce because in such a virtual environment, participants are usually anonymous and do not engage in direct face-to-face communication (Cofta 2006, Kim et al. 2009, Lu et al. 2010, Pentina et al. 2013). Moreover, the growing number of fraudulent practices has discouraged customers from adopting e-commerce (Lek et al. 2001). On the other hand, enhancing the degree of trust can increase online purchase intention (Gefen 2000) and help maintain long-term relationships between businesses and customers (Hoffman et al. 1999, Reichheld & Schefter 2000).

Since its collaboration with Taobao, Sina Weibo grants the verified Taobao sellers

an identity of "Tao," which is highlighted in their Sina Weibo profiles. Furthermore, like Twitter, Sina Weibo also supports verification for individuals or entities[4]. After being approved, the verified user will have a "V"[5] identity in the profile. Both the "Tao" and "V" identities provide a signal of trust and customers are more inclined to interact with the verified sellers. Since the verification of identities is associated with some thresholds of credibility and prestige[6], it can also be a proxy of other unmeasurable factors such as the quality of service. Hence, the sellers with verified identities are more likely to gain popularity in social media.

As discovered in previous studies, we expect that the number of followers has a positive effect on a seller's marketing popularity (Wang & Jin 2010), so as the average number of pictures posted per marketing microblog (De Vries et al. 2012, Fortin & Dholakia 2005). Some studies have also suggested the gender differences in online activity on Sina Weibo (Guan et al. 2014, Li et al. 2015) and we also control the dummy variable of gender in our analysis.

3 Data and Methods

In this section we provide detailed information on how we gather the data from Sina Weibo. We then clean the data and, most importantly, classify the microblogs

[4]There are a number of requirements such as the number of followers being at least 100 and the number of followees being at least 30, see `http://verified.weibo.com/verify/help?fr=home&frpos=leftnav` for more details.

[5]Sina Weibo accepts voluntary requests from the elite of 34 categories and 542 professions which include electronic retailers (see `http://verified.weibo.com/verify/applystd?fr=home&frpos=morestd` for more details).

[6]See footnotes 4 and 5 above.

to construct the key variable of marketing aggressiveness. Finally, we discuss the modelling methods and results of testing our hypothesis on the relationship between the marketing aggressiveness level and the marketing popularity.

3.1 Data Sampling

First, we identify the Taobao sellers on the microblog platform Sina Weibo in November 2014 using data scraping techniques. We select this period because November 11 (a.k.a. Singles' Day[7]), the largest online shopping day in the world[8], occurs in this month. As a result, the Taobao sellers should have substantial economic incentives to post marketing microblogs during this period. Taking advantage of the search function of Sina Weibo, we find 281,160 profiles including the Chinese characters "Taobao" in their personal labels. However, apart from the 19,309 users with a verified "Tao" badge, we cannot make sure whether all the other users are real Taobao sellers or just fond of shopping on Taobao website. Hence, we retain only the users who add a link of Taobao shop or at least one of the Chinese words meaning Taobao seller[9] in their personal tags. Finally we identify 52,187 Taobao sellers and collect the 465,812 microblogs they posted in November 2014. Considering that the sellers may be more aggressive in promoting their products in the period of Singles' Day and this may introduce biases to our sample, we further track the marketing behavior of

[7]Singles' Day is a day for people who are single, celebrated on November 11 (11/11). The date is chosen for the connection between singles and the number "1." This holiday has become popular in recent years among young Chinese people.

[8]The sales of Alibaba's sites Tmall and Taobao are $9.3 billion on November 11th, 2014.

[9]See some examples in Figure A2.

12,744 sellers who add the links of their Taobao shops[10] in normal months from July to October in 2016 and collect the 308,167 microblogs they posted.

3.2 Classification of Microblogs

A traditional text classification framework comprises preprocessing, feature extraction, feature selection and classification steps (Allahyari et al. 2017). In the preprocessing step, we first randomly select 5,000 out of the 774,429 microblogs and manually label them as either "marketing" or "non-marketing." We employ these microblogs to train the classifiers so as to predict the labels of all the microblogs. The second process is tokenization, which is a task of breaking a character sequence and a defined document unit into pieces such as words, phrases, symbols and other elements called tokens (Manning & Schutze. 2008). Using the package "*jieba*" in Python which is designed for Chinese words segmentation, we split each of the 5,000 microblogs into a list of words.

Most machine learning algorithms require numerical feature vectors with a fixed size rather than the raw texts with variable length. Hence, we need to extract numerical features from the contents instead of using the symbols directly. An intuitive way is to assign a weight to each word in a given document. We here adopt the method of *term frequency - inverse document frequency* (TF-IDF), which is a numerical statistic that is intended to reflect how important a word is to a document in a corpus (Salton & Buckley 1988, Leskovec et al. 2014). In this step we filter out some

[10]By adding the links, they are classified as more dedicated social media marketers and are more likely to remain active in later periods.

stop words[11] and the words that appear only once. In our case, each row in the TF-IDF matrix $\mathbf{A}$ represents a microblog d and each column corresponds to a word t. The term frequency $tf(d,t)$ is the number of times that word t appears in the microblog d. The inverse document frequency $idf(t)$ is equal to $\log(\frac{N+1}{n_t+1})+1$, where N is the number of microblogs and n_t is the number of microblogs containing word t. The corresponding value for word t in microblog d in the matrix $\mathbf{A}$ is defined as $\mathbf{A}(t,d) = tf(d,t) * idf(t)$. Table 1 presents several rows and columns of the TF-IDF matrix of our sample. The TF-IDF matrix will be used as inputs to predict the labels of the microblogs in the next procedures.

Table 1: Term Frequency-Inverse Document Frequency Matrix

	Terms	*http*	11	*purchasing agent*	*really*	*new fashion*	⋯
label	Documents						
0	microblog 1	0.173	0	0	0	0	⋯
1	microblog 2	0.069	0.182	0	0	0.201	⋯
1	microblog 3	0	0	0.192	0	0	⋯
⋯	⋯			⋯			
0	microblog 5000	0	0	0	0	0	⋯

Notes: we label the "marketing" microblogs as 1 and "non-marketing" ones as 0.

Given that the classification task in our case belongs to supervised learning, we perform the following steps. First, we choose four machine learning algorithms to classify separately the microblogs, including logistic regression (logit), decision tree[12] (CART), random forest (RndFor) and multinomial naive bayes (MNB)[13]. The first two algorithms require that only few features can be introduced into the model.

[11]The stop words refer to some extremely common words that would appear to be of little value in helping select documents matching a user's need (Manning & Schutze. 2008).

[12]We use the GINI criterion in the splitting decision.

[13]See the details of the four algorithms in Baesens (2014) and Baesens et al. (2015).

To reduce the dimension, we perform the truncated singular vector decomposition (TSVD) method (Manning & Schutze. 2008). Then we select the main features extracted by TSVD in the two algorithms.[14] When using RndFor and MNB algorithms, we directly use the TF-IDF matrix as inputs.

Second, we randomly select 4,000 out of the 5,000 labelled microblogs as the training set and the rest 1,000 microblogs as the test set. Since the algorithms logit, CART and RndFor require us to determine the parameters, we apply a cross-validation method[15] to optimize them. Using a 10-fold cross-validation approach, we randomly divide the training set into 10 groups of equal size. Each time we use a given algorithm to fit 9 folds and evaluate its performance on the rest 1 fold, which is the validation set. The evaluation is based on the *area under the receiver operating characteristics curve* (AUC) score, which is the most informative and objective indicator of predictive accuracy within a benchmarking context (Lessmann et al. 2008). This procedure is repeated 10 times and each time a different group of observations is treated as the validation set. For each set of parameters, we compute the average of the 10 AUC scores. Then we compare the average AUC score of different parameter settings and select the best performing set of parameters. We adopt the parameters' values suggested in Lessmann et al. (2008)[16], and the ones with the best average

[14]In logit, we further apply the recursive feature selection to remove several less important features.

[15]See the details in James et al. (2014).

[16]The parameters' values suggested in Lessmann et al. (2008) for each classifier are as follows: for logit and CART, we first perform dimension reduction to extract the k main components where k varies from 5 to 20. The 5 main components explain 23.0% of the variance while the 20 main components explain 28.9%. In logit we further perform the recursive feature elimination to select the inputs in the regression model and the number of features selected varies from 5 to k. As for RndFor, the parameter to be tuned is the number of trees, and the suggested values include [10,

AUC score of each algorithm are listed in Table 2.

Table 2: Hyperparameters with Optimal Test Performance

	num_components	*num_features_selected*	*num_trees*	*AUC*
logit	19	14		0.854
CART	20	20		0.765
RndFor			250	0.956
MNB				0.960

Table 3: Performance Matrix

	precision rate	*recall rate*	*accuracy rate*	$f1$ *score*	$confusion\ matrix(TP, FP, FN, TN)$
logit	0.792	0.455	0.772	0.578	[156,187,41,616]
CART	0.645	0.700	0.765	0.671	[240,103,132,525]
RndFor	0.888	0.810	0.900	0.848	[278,65,35,622]
MNB	0.913	0.770	0.896	0.835	[264,79,25,632]

Finally, after we determine the parameters for each algorithm, we assess their performance on the test set and report the results in Table 3. To further understand which algorithm has a statistically significant better performance, we conduct the test as in DeLong et al. (1988). From Table 4 we observe that RndFor and MNB are significantly better than logit and CART. But there is no significant difference between RndFor and MNB. Since MNB has the largest AUC value on the test set, we adopt it to classify all the rest microblogs.

3.3 Modelling Method

After we classify all the microblogs, we aggregate the data by seller and thus obtain the number of marketing-related microblogs each seller posts in each observation

50, 100, 250, 500, 1,000].

Table 4: DeLong's test to compare different models' AUC scores

	AUC	logit	CART	RndFor	MNB
logit	0.854	-			
CART	0.765	0.000	-		
RndFor	0.956	0.000	0.000	-	
MNB	0.960	0.000	0.000	0.496	-

*Except the AUC scores, the above listed numbers are p-values.

period. Then we compute the proportion of the number of marketing microblogs to the total number of microblogs, which is used to measure marketing aggressiveness. We also count the average number of likes per marketing microblog for each seller, which is used to measure marketing popularity. The sellers whose average number of likes located in the 1% tail of the distribution are excluded from the sample in order to reduce the possible outlier effect. Furthermore, we remove the sellers who posted less than 10 microblogs in each observation period since the possible wrong classification may cause a large error on the value of marketing proportion for sellers who barely post. The details of the data in each of the previous steps are presented in Table A1.

In the final sample, the response variable is the average number of likes a seller receives per marketing microblog ($average\_likes\_marketing$). The inputs include the marketing aggressiveness level ($marketing\_proportion$), the number of followers ($followers$), the gender ($female$), and whether the seller has the "V" (V) and the "Tao" (Tao) identities. For the data of 2016, we manage to extract the average number of pictures a seller posts in each marketing-related microblog (num_pic) as well. The definitions of the variables are shown in Table 5 and their descriptive

statistics are shown in Table 6. On average, a seller receives correspondingly 0.604 and 0.641 likes per marketing microblog in 2014 and 2016. The standard deviations of the variable in the two data sets are also quite close, which are respectively 1.107 and 1.150. The average proportion of marketing microblogs is 33.3% for the data of 2014 and 28.5% for 2016.

Table 5: Variable Definitions

Variables	Definitions
average_likes_marketing	The average number of likes received from all the marketing-related microblogs.
marketing_proportion	The proportion of the number of marketing-related microblogs to the total number of microblogs.
Tao	1 if the seller has a verified "Tao" identity; otherwise 0.
V	1 if the seller's real identity has been verified; otherwise 0.
female	1 if the seller is a female; otherwise 0.
followers	The number of followers the seller has.
num_pic	The average number of pictures the seller posts per marketing-related microblog.

Table 6: Summary Statistics

	2014					2016				
Variables	N	mean	sd	min	max	N	mean	sd	min	max
average_likes_marketing	5,809	0.604	1.107	0	7.25	1,812	0.641	1.150	0	5
marketing_proportion	5,809	0.333	0.302	0.0034	1	1,812	0.285	0.292	0.0009	1
female	5,809	0.759	0.428	0	1	1,812	0.829	0.376	0	1
followers	5,809	9,830	46,553	1	1,011,968	1,812	14,613	83,851	4	1,866,129
Tao	5,809	0.583	0.493	0	1	1,812	0.512	0.500	0	1
V	5,809	0.209	0.407	0	1	1,812	0.221	0.415	0	1
num_pic						1,812	3.804	2.790	0	9

Next we perform linear regression, multi-layer perceptron neural network (MLP) and random forest to explore the relationship between *marketing_proportion* and *average_likes_marketing*. In the linear regression, we add the squared term of *marketing_proportion* to test our hypothesis. Moreover, since some inputs may have a non-linear effect on the response variable, we follow the steps introduced in Van Gestel et al. (2006) and Van Gestel et al. (2005) to perform the Yeo Johnson

transformation (Yeo & Johnson 2000) for the continuous variables in the linear regression (see details in Appendix A). In MLP and RndFor, we use 5-fold cross-validation to select the hyperparameters with the best performance[17].

4 Results and Discussion

Two-thirds of each data set of the two years are randomly selected to train the algorithms and the rest one-third are used to test their performance. We evaluate their performance based on three indicators: R-square, mean square error (MSE) and Pearson correlation coefficient between the predicted value and the true value. As Table 7 shows, the linear regression using Yeo Johnson transformation[18] has the best performance compared with the standard linear regression, MLP[19] and RndFor[20].

Table 8 presents the estimation results of the linear regression using Yeo-Johnson Transformation. We notice that the coefficients of *marketing_proportion* and its quadratic term are both significant, and as expected, the coefficients of the quadratic term are negative for both data sets. To confirm that the relationship follows an inverted U-shape, we further perform the test in Lind & Mehlum (2010) (see details

[17]Lessmann et al. (2008) assume that there is a single hidden layer for MLP and the hyperparmeters' values of the number of neurons are [4,5,6]. As for RndFor, the number of trees are seleced from [10, 50, 100, 250, 500, 1,000].

[18]We perform the Yeo Johnson transformation for all the continuous variables, but only the transformation of the number of followers significantly improves the performance of the model. The estimated hyperparameters for the number of followers are $\lambda = -2, c = -0.5$ for the data of 2014 and $\lambda = -2, c = 0$ for data of 2016.

[19]In MLP, the optimal number of neurons is 4 for both the data of 2014 and 2016.

[20]In RndFor, the optimal number of trees is 250 for both the data of 2014 and 1000 for 2016.

Table 7: Performance of Different Models on Test Set

	2014			2016		
Models	R^2	MSE	Pearson	R^2	MSE	Pearson
Linear Regression	0.081	1.085	0.298	0.220	1.039	0.469
Yeo Johnsson	0.247	0.889	0.498	0.369	0.841	0.608
MLP	0.038	1.136	0.230	0.111	1.184	0.341
RndFor	0.134	1.023	0.421	0.308	0.922	0.560

Table 8: Regression Results of the Yeo-Johnson Transformation (Training Set)

VARIABLES	2014	2016
marketing_proportion	0.567***	1.185***
	(0.189)	(0.315)
marketing_proportion2	-1.021***	-1.623***
	(0.183)	(0.313)
V	0.0157	0.0133
	(0.0483)	(0.0755)
Tao	0.0271	-0.0786
	(0.0324)	(0.0546)
transform_followers	1.296***	6.204***
	(0.0715)	(0.404)
female	0.272***	0.0310
	(0.0357)	(0.0640)
num_pic		0.0418***
		(0.0105)
Constant	1.488***	0.159**
	(0.0851)	(0.0762)
Observations	3,872	1,208
R-squared	0.241	0.359
Adjusted R-squared	0.240	0.355

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

in Appendix). Table 9 shows that, in both data sets, the slope at the lower bound of the data range is significant and positive and the slope at the upper bound is significant and negative. Moreover, the turning points are respectively 0.278 and 0.365, and their 95% confidence intervals are respectively [0.146, 0.346] and [0.272, 0.416] (see details in Appendix B), which are located in the data range (0, 1]. Thus, the above results strongly support our hypothesis that there is an inverted U-shape relationship between *average_likes_marketing* and *marketing_proportion*. Figure 1 and Figure 2 visualise the inverted U-shape relationship[21]. The results imply that when the proportion of marketing-related microblogs increases under a certain level, people become more likely to respond with liking their microblogs. However, if the proportion continues to increase beyond the certain level, the trend reverts and the social media marketers get penalised for being too aggressive.

Table 9: Inverted U-shape Test

	2014		2016	
	Lower bound	Upper bound	Lower bound	Upper bound
Interval	0.003	1	0.001	1
Slope	0.560	-1.474	1.182	-2.060
t-value	2.985	-7.939	3.761	-6.304
$P > \|t\|$	0.001	0.000	0.000	0.000
Overall test	t-value: 2.98		t-value: 3.76	
	$P > \|t\|$: 0.001		$P > \|t\|$: 0.000	
Turning point	0.278		0.365	
95% CI	[0.146, 0.346]		[0.272, 0.416]	

[21]In the graphs, the average numer of likes are computed given the mean value of the transformed number of followers and the number of pictures, and female equal to 1 for the data of 2014. The variables that are not significant are ignored in the calculation.

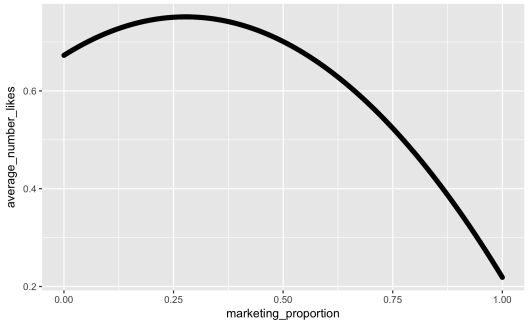

Figure 1: Optimal Proportion 2014

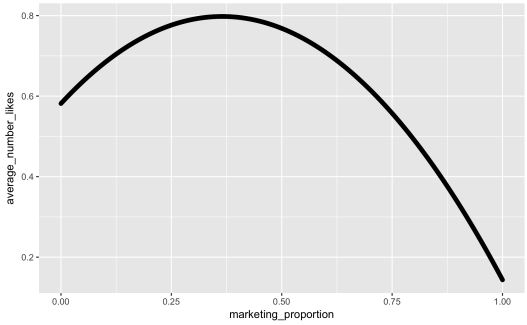

Figure 2: Optimal Proportion 2016

We make some further robustness checks by performing other specifications such as using the cubic, logarithm and exponential forms of the *marketing_proportion*. As Table A2 shows, the performance of these specifications are almost the same. Though we cannot exclude other possibilities, what we empirically find is always that when the proportion of marketing-related microblogs goes beyond a certain level, sellers receive fewer likes from their followers.

Interestingly, we also notice that the non-linear transformation of the number of followers significantly improves the prediction accuracy of the number of likes. The relationship between the number of likes and followers of the two years data is depicted in Figure 3 and Figure 4. It can be seen that an increase of 10,000 followers from 10,000 to 20,000 has a much larger effect on the average number of likes than an increase from 100,000 to 110,000. That is, after the number of followers reaches around 100,000, a saturation effect occurs.

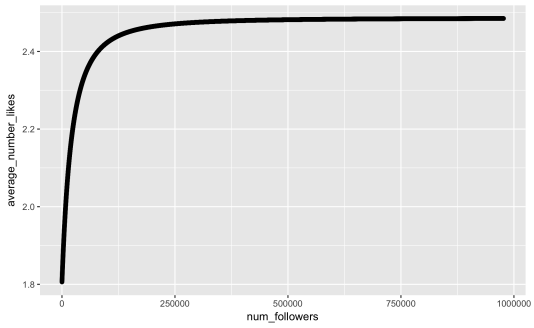

Figure 3: Likes and Followers 2014

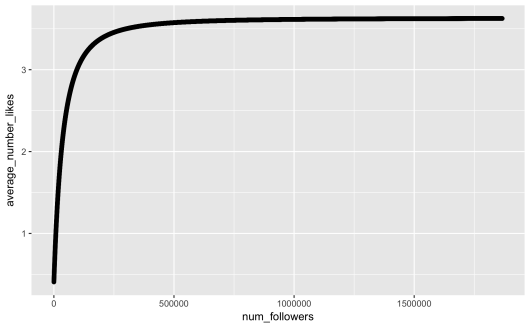

Figure 4: Likes and Followers 2016

5 Conclusion

In this paper, we empirically investigate the relationship between sellers' marketing aggressiveness level and their marketing popularity in social media. In particular, by analysing the microblogs of Taobao sellers on Sina Weibo, we find that the relationship between the proportion of marketing-related microblogs (our measure of marketing aggressiveness) and the average number of likes a seller receives per marketing microblog (our measure of marketing popularity) follows an inverted U-shaped curve. The corresponding optimal proportion of marketing microblogs is around 0.3.

Our paper contributes to the literature of social media marketing by highlighting the tension between customers' social and commercial needs in the C2C context. Our finding suggests a good practice of social media marketing to strike a balance between the two. That is, a seller may get penalised with less popularity for being too aggressive and hence crowding out customers' social content in the limited social media space. Our paper also contributes to the recent literature of social media metrics (Peters et al. 2013) and our measure of marketing aggressiveness forms an

informative metric for management and marketing practitioners to monitor in real time and to improve their decision making process. The marketing input can be adjusted accordingly when the aggressiveness level becomes too high or too low. Creative marketing microblogs can also be composed to meet customers' social needs as well as to serve the commercial purpose.

To extend the current study, we plan to investigate the impact of social media marketing on the C2C sellers' sales by collecting data from both Taobao and Sina Weibo. We would like to see if the linkage can be established between the marketing popularity in Sina Weibo and the ultimate marketing outcome, sales, in Taobao. Another opportunity after having the cross-platform data is to explore how the sales performance in Taobao affects the marketing activities such as the aggressiveness level in Sina Weibo. Our paper also suggests to examine the possibility of combining the social and commercial contents in a single marketing microblog. The current study employs a binary classification for the microblogs. We therefore leave it to future research to identify the third type and to quantify how effective it is in terms of attaining popularity.

References

Akar, E. & Topçu, B. (2011), 'An examination of the factors influencing consumers' attitudes toward social media marketing', *Journal of Internet Commerce* **10**(1), 35–67.

Allahyari, M., Pouriyeh, S., Assefi, M., Safaei, S., Trippe, E. D., Gutierrez, J. B. &

Kochut, K. (2017), 'A brief survey of text mining: Classification, clustering and extraction techniques', *arXiv preprint arXiv:1707.02919* .

Baesens, B. (2014), *Analytics in a big data world: The essential guide to data science and its applications*, John Wiley & Sons.

Baesens, B., Van Vlasselaer, V. & Verbeke, W. (2015), *Fraud analytics using descriptive, predictive, and social network techniques: a guide to data science for fraud detection*, John Wiley & Sons.

Ballings, M., Van den Poel, D. & Bogaert, M. (2016), 'Social media optimization: Identifying an optimal strategy for increasing network size on facebook', *Omega* **59**, 15–25.

Bauer, H. H., Reichardt, T., Barnes, S. J. & Neumann, M. M. (2005), 'Driving consumer acceptance of mobile marketing: A theoretical framework and empirical study', *Journal of electronic commerce research* **6**(3), 181.

Chu, S.-C. & Choi, S. M. (2011), 'Electronic word-of-mouth in social networking sites: a cross-cultural study of the united states and china', *Journal of Global Marketing* **24**(3), 263–281.

Chu, S.-C. & Kim, Y. (2011), 'Determinants of consumer engagement in electronic word-of-mouth (ewom) in social networking sites', *International journal of Advertising* **30**(1), 47–75.

Cofta, P. (2006), 'Convergence and trust in ecommerce', *BT Technology Journal* **24**(2), 214–218.

De Vries, L., Gensler, S. & Leeflang, P. S. (2012), 'Popularity of brand posts on brand fan pages: An investigation of the effects of social media marketing', *Journal of Interactive Marketing* **26**(2), 83–91.

DeLong, E. R., DeLong, D. M. & Clarke-Pearson, D. L. (1988), 'Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach', *Biometrics* pp. 837–845.

Eighmey, J. & McCord, L. (1998), 'Adding value in the information age: Uses and gratifications of sites on the world wide web', *Journal of business research* **41**(3), 187–194.

Fieller, E. C. (1954), 'Some problems in interval estimation', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 175–185.

Fortin, D. R. & Dholakia, R. R. (2005), 'Interactivity and vividness effects on social presence and involvement with a web-based advertisement', *Journal of business research* **58**(3), 387–396.

Gangadharbatla, H. (2008), 'Facebook me: Collective self-esteem, need to belong, and internet self-efficacy as predictors of the igenerations attitudes toward social networking sites', *Journal of interactive advertising* **8**(2), 5–15.

Gefen, D. (2000), 'E-commerce: the role of familiarity and trust', *Omega* **28**(6), 725–737.

Gefen, D. & Straub, D. W. (2004), 'Consumer trust in b2c e-commerce and the

importance of social presence: experiments in e-products and e-services', *Omega* **32**(6), 407–424.

Grant, I. C. (2005), 'Young peoples' relationships with online marketing practices: An intrusion too far?', *Journal of Marketing Management* **21**(5-6), 607–623.

Guan, W., Gao, H., Yang, M., Li, Y., Ma, H., Qian, W., Cao, Z. & Yang, X. (2014), 'Analyzing user behavior of the micro-blogging website sina weibo during hot social events', *Physica A: Statistical Mechanics and its Applications* **395**, 340–351.

Haans, R. F., Pieters, C. & He, Z.-L. (2015), 'Thinking about u: theorizing and testing u-and inverted u-shaped relationships in strategy research', *Strategic Management Journal* .

He, Y. & Song, L. (2015), 'The study on factors affecting the attraction of microblog', *Computers in Human Behavior* .

Hoffman, D. L., Novak, T. P. & Peralta, M. (1999), 'Building consumer trust online', *Communications of the ACM* **42**(4), 80–85.

James, G., Witten, D. & Hastie, T. (2014), 'An introduction to statistical learning: With applications in r.'.

Kaplan, A. M. & Haenlein, M. (2010), 'Users of the world, unite! the challenges and opportunities of social media', *Business horizons* **53**(1), 59–68.

Katz, E., Haas, H. & Gurevitch, M. (1973), 'On the use of the mass media for important things', *American sociological review* pp. 164–181.

Kim, H.-b., Kim, T. T. & Shin, S. W. (2009), 'Modeling roles of subjective norms and etrust in customers' acceptance of airline b2c ecommerce websites', *Tourism management* **30**(2), 266–277.

King, R. A., Racherla, P. & Bush, V. D. (2014), 'What we know and don't know about online word-of-mouth: A review and synthesis of the literature', *Journal of Interactive Marketing* **28**(3), 167–183.

Kumar, V. & Mirchandani, R. (2012), 'Increasing the roi of social media marketing', *MIT sloan management review* **54**(1), 55.

Lek, M., Anadarajah, B., Cerpa, N. & Jamieson, R. (2001), 'Data mining prototype for detecting ecommerce fraud', *ECIS 2001 Proceedings* p. 60.

Leskovec, J., Rajaraman, A. & Ullman, J. D. (2014), *Mining of massive datasets*, Cambridge university press.

Lessmann, S., Baesens, B., Mues, C. & Pietsch, S. (2008), 'Benchmarking classification models for software defect prediction: A proposed framework and novel findings', *IEEE Transactions on Software Engineering* **34**(4), 485–496.

Li, Y., Gao, H., Yang, M., Guan, W., Ma, H., Qian, W., Cao, Z. & Yang, X. (2015), 'What are chinese talking about in hot weibos?', *Physica A: Statistical Mechanics and its Applications* **419**, 546–557.

Lin, K.-Y. & Lu, H.-P. (2011), 'Why people use social networking sites: An empirical study integrating network externalities and motivation theory', *Computers in Human Behavior* **27**(3), 1152–1161.

Lind, J. T. & Mehlum, H. (2010), 'With or without u? the appropriate test for a u-shaped relationship', *Oxford bulletin of economics and statistics* **72**(1), 109–118.

Lu, Y., Zhao, L. & Wang, B. (2010), 'From virtual community members to c2c e-commerce buyers: Trust in virtual communities and its effect on consumers purchase intention', *Electronic Commerce Research and Applications* **9**(4), 346–360.

Mangold, W. G. & Faulds, D. J. (2009), 'Social media: The new hybrid element of the promotion mix', *Business horizons* **52**(4), 357–365.

Manning, Christopher, P. R. & Schutze., H. (2008), *Introduction to information retrieval.*, Cambridge University Press.

Michaelidou, N., Siamagka, N. T. & Christodoulides, G. (2011), 'Usage, barriers and measurement of social media marketing: An exploratory investigation of small and medium b2b brands', *Industrial Marketing Management* **40**(7), 1153–1159.

Pentina, I., Zhang, L. & Basmanova, O. (2013), 'Antecedents and consequences of trust in a social media brand: A cross-cultural study of twitter', *Computers in Human Behavior* **29**(4), 1546–1555.

Peters, K., Chen, Y., Kaplan, A. M., Ognibeni, B. & Pauwels, K. (2013), 'Social media metricsa framework and guidelines for managing social media', *Journal of interactive marketing* **27**(4), 281–298.

Rapp, A., Beitelspacher, L. S., Grewal, D. & Hughes, D. E. (2013), 'Understanding

social media effects across seller, retailer, and consumer interactions', *Journal of the Academy of Marketing Science* **41**(5), 547–566.

Reichheld, F. F. & Schefter, P. (2000), 'E-loyalty', *Harvard business review* **78**(4), 105–113.

Ruggiero, T. E. (2000), 'Uses and gratifications theory in the 21st century', *Mass communication & society* **3**(1), 3–37.

Salton, G. & Buckley, C. (1988), 'Term-weighting approaches in automatic text retrieval', *Information processing & management* **24**(5), 513–523.

Sasabuchi, S. (1980), 'A test of a multivariate normal mean with composite hypotheses determined by linear inequalities', *Biometrika* **67**(2), 429–439.

Schivinski, B. & Dabrowski, D. (2014), 'The effect of social media communication on consumer perceptions of brands', *Journal of Marketing Communications* pp. 1–26.

Swani, K., Brown, B. P. & Milne, G. R. (2014), 'Should tweets differ for b2b and b2c? an analysis of fortune 500 companies' twitter communications', *Industrial Marketing Management* **43**(5), 873–881.

Taylor, D. G., Lewin, J. E. & Strutton, D. (2011), 'Friends, fans, and followers: Do ads work on social networks?.', *Journal of Advertising Research* **51**(1), 258 – 275.

Teo, T. S. & Liu, J. (2007), 'Consumer trust in e-commerce in the united states, singapore and china', *Omega* **35**(1), 22–38.

Tsai, W.-H. S. & Men, L. R. (2013), 'Motivations and antecedents of consumer

engagement with brand pages on social networking sites', *Journal of Interactive Advertising* **13**(2), 76–87.

Van Gestel, T., Baesens, B., Van Dijcke, P., Garcia, J., Suykens, J. A. & Vanthienen, J. (2006), 'A process model to develop an internal rating system: Sovereign credit ratings', *Decision Support Systems* **42**(2), 1131–1151.

Van Gestel, T., Baesens, B., Van Dijcke, P., Suykens, J., Garcia, J. & Alderweireld, T. (2005), 'Linear and nonlinear credit scoring by combining logistic regression and support vector machines', *Journal of credit Risk* **1**(4).

Wang, R. & Jin, Y. (2010), An empirical study on the relationship between the followers' number and influence of microblogging, *in* 'E-Business and E-Government (ICEE), 2010 International Conference on', IEEE, pp. 2014–2017.

Wiersema, F. (2013), 'The b2b agenda: The current state of b2b marketing and a look ahead', *Industrial Marketing Management* **4**(42), 470–488.

Yeo, I.-K. & Johnson, R. A. (2000), 'A new family of power transformations to improve normality or symmetry', *Biometrika* pp. 954–959.

Yu, L., Asur, S. & Huberman, B. A. (2011), 'What trends in chinese social media', *arXiv preprint arXiv:1107.3522* .

Appendix

Figure A1: An Example of a Taobao Seller's Sina Weibo Account

淘宝卖家	淘宝店主	淘宝掌柜	淘宝商户

Figure A2: Variants of "Taobao Sellers" in Chinese

Table A1: Statistics by steps

Data of 2014				
Steps	num_sellers	num_microblogs	num_marketing_microblogs	proportion of marketing microblogs
1. Initial sample	52,187	465,812	152,772	32.8%
2. sellers who post	18,809	465,812	152,772	32.8%
3. sellers who market	9,577	402,533	152,772	38.0%
4. final sample*	5,809	385,911	145,328	37.7%

Data of 2016				
Steps	num_sellers	num_microblogs	num_marketing_microblogs	proportion of marketing microblogs
1. Initial sample	12,744	308,617	86,115	27.9%
2. sellers who post	4,227	308,617	86,115	27.9%
3. sellers who market	2,262	217,806	86,115	39.5%
4. final sample*	1,812	212,717	84,724	39.8%

* In the final sample we remove the outliers and sellers who post less than 10 microblogs.
The proportion of marketing microblogs in the last row is computed by $\frac{\sum_i No.marketing\ microblogs_i}{\sum_i No.total\ microblogs_i}$ while the one in Table 6 is computed by $\frac{\sum_i marketing\ proportion_i}{n}$.

Appendix A: Estimation of Yeo-Johnson Transformation

Yeo & Johnson (2000) has proposed a transformation which is of the same form as Box-Cox transformations and is also valid for negative values. The transformation function is defined as follows:

$$f(x;\lambda) = \begin{cases} ((1+x)^{\lambda}-1)/\lambda, & \lambda \neq 0, x \geq 0, \\ \log(x+1), & \lambda = 0, x \geq 0, \\ -((1-x)^{(2-\lambda)}-1)/(2-\lambda), & \lambda \neq 2, x \leq 0, \\ -\log(-x+1), & \lambda = 2, x \leq 0. \end{cases}$$

We consider the following transformation: $x \mapsto f(x+c,\lambda)$ where c is the location parameter and λ is the transformation parameter. The parameters c and λ are esti-

mated based on the following steps:

Step 1: The variable to be transformed is first normalized to zero median and unit variance.

Step 2: The parameters are estimated using a grid search mechanism. The parameter c varies from -3 to $+3$ and the parameter λ varies from -2 to $+2$. For each hyperparameter combination (c, λ), the model is estimated and MSE is stored.

Step 3: Using a 5-fold cross-validation, the combination (c, λ) with the lowest average MSE is selected. The optimal MSE is compared with the MSE obtained with $\lambda = 1$. When the MSE of the nonlinear model is lower than the MSE of the linear model, the nonlinear transformation is applied.

Appendix B: Test of the Inverted U-shape

$$Y = \beta_0 + \beta_1 X + \beta_2 X^2 + \boldsymbol{\gamma}' \boldsymbol{Z} + \epsilon$$

Haans et al. (2015) framed the U-shape test proposed in Lind & Mehlum (2010) as a three-step procedure. Here we list the steps to test whether the relationship between Y and X is an inverted U-shape:

Step 1: β_2 needs to be significant and negative.

Step 2: The slope must be significantly steep at both ends of the data range $[X_l, X_u]$, where X_l is the minimal of X and X_u is the maximal. To make sure that the inverted U-shape is a phenomenon in the interior of the range of X, the slope at the lower bound $\beta_1 + 2\beta_2 X_l$ should be significant and positive and the slope at the upper bound $\beta_1 + 2\beta_2 X_u$ should be significant and negative. Hence, we need to test whether the combined null hypothesis can be rejected in favour of the alternative hypothesis:

$H_0 : \beta_1 + 2\beta_2 X_l \leq 0$ and/or $\beta_1 + 2\beta_2 X_u \geq 0$

$H_1 : \beta_1 + 2\beta_2 X_l > 0$ and $\beta_1 + 2\beta_2 X_u < 0$

Sasabuchi (1980) provides a test based on the likelihood ratio principle. The rejection areas are as follows:

$$R_\alpha = \{(\beta_1, \beta_2) : \frac{\beta_1 + 2\beta_2 X_l}{\sqrt{s_{11} + 4X_l s_{12} + 4X_l^2 s_{22}}} > t_\alpha \ \& \ \frac{\beta_1 + \beta_2 2X_u}{\sqrt{s_{11} + 4X_u s_{12} + 4X_u^2 s_{22}}} < -t_\alpha\}$$

where s_{11}, s_{22}, s_{12} are the estimated variance of β_1 and β_2 and their covariance, and t_α is the α percentile of the t-distribution with the appropriate degree of freedom.

Step 3: The turning point and its 95% confidence interval needs to be located within

the range of X. The point estimate of the turning point is $X = -\frac{\beta_2}{2\beta_1}$. Fieller (1954) provides how to construct a confidence interval for the ratio of two normally distributed estimates. The lower bound and upper bound of the $(1 - 2\alpha)$ confidence interval for $-\frac{\beta_2}{2\beta_1}$ are respectively

$$\widetilde{X}_l = \frac{s_{12}t_\alpha^2 - \beta_1\beta_2 - t_\alpha\sqrt{(s_{12}^2 - s_{22}s_{11})t_\alpha^2 + \beta_2^2 s_{11} + \beta_1^2 s_{22} - 2s_{12}\beta_1\beta_2}}{\beta_2^2 - s_{22}t_\alpha^2}$$

$$\widetilde{X}_u = \frac{s_{12}t_\alpha^2 - \beta_1\beta_2 + t_\alpha\sqrt{(s_{12}^2 - s_{22}s_{11})t_\alpha^2 + \beta_2^2 s_{11} + \beta_1^2 s_{22} - 2s_{12}\beta_1\beta_2}}{\beta_2^2 - s_{22}t_\alpha^2}$$

If the confidence interval is located within the range of X, we can make sure that the relationship between Y and X is an inverted U-shape.

Appendix C: Other Specifications of Linear Regression

Table A2: Other Specifications of Linear Regression

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	2014_cubic	2014_log	2014_exp	2016_cubic	2016_log	2016_exp
marketing_proportion	2.566***			2.026***		
	(0.481)			(0.723)		
marketing_proportion2	-6.538***			-4.052**		
	(1.152)			(1.816)		
marketing_proportion3	3.819***			1.710		
	(0.746)			(1.202)		
log_marketing_proportion		-0.0456***			-0.00848	
		(0.0132)			(0.0178)	
exp_marketing_proportion			-0.247***			-0.203***
			(0.0251)			(0.0444)
V	0.0166	0.0298	0.0185	0.0178	0.0130	0.00742
	(0.0482)	(0.0489)	(0.0485)	(0.0752)	(0.0765)	(0.0761)
Tao	0.0224	0.0204	0.0369	-0.0795	-0.0866	-0.0650
	(0.0322)	(0.0325)	(0.0323)	(0.0546)	(0.0551)	(0.0550)
transform_followers	1.294***	1.316***	1.308***	6.169***	6.367***	6.427***
	(0.0712)	(0.0724)	(0.0717)	(0.407)	(0.409)	(0.405)
female	0.267***	0.275***	0.286***	0.0344	0.0301	0.0288
	(0.0355)	(0.0358)	(0.0358)	(0.0643)	(0.0640)	(0.0639)
num_pic				0.0416***	0.0427***	0.0436***
				(0.0105)	(0.0107)	(0.0107)
Constant	1.365***	1.410***	1.825***	0.112	0.204***	0.488***
	(0.0891)	(0.0818)	(0.0902)	(0.0864)	(0.0737)	(0.0902)
Observations	3,872	3,872	3,872	1,208	1,208	1,208
R-squared	0.245	0.228	0.238	0.360	0.343	0.350

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

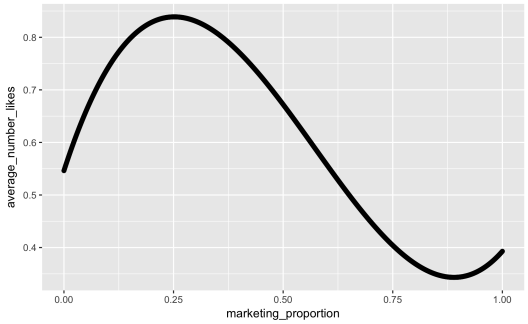

Figure A3: Cubic Specification 2014

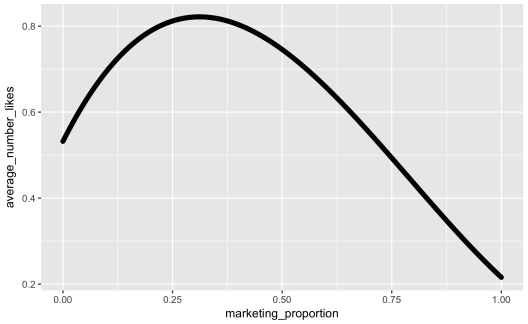

Figure A4: Cubic Specification 2016

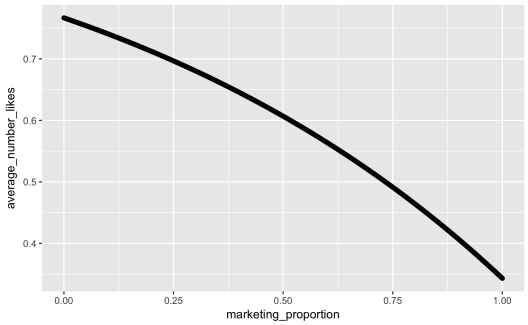

Figure A5: Exp. Specification 2014

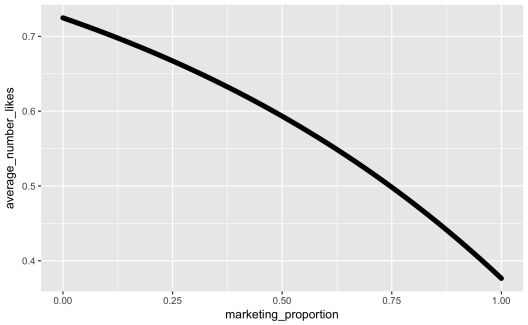

Figure A6: Exp. Specification 2016