2-9-2016

What Guests Really Think of Your Hotel: Text Analytics of Online Customer Reviews

Hyun Jeong "Spring" Han
National University Higher School of Economics in Moscow

Shawn Mankad
Cornell University

Nagesh Gavirneni
Cornell University

Rohit Verma
Cornell University, rv54@cornell.edu

Recommended Citation

Han, H. J., Mankad, S., Gavirneni, N., & Verma, R. (2016). What guests really think of your hotel: Text analytics of online customer reviews. *Cornell Hospitality Report, 16*(2), 3-17.

What Guests Really Think of Your Hotel: Text Analytics of Online Customer Reviews

Abstract

Hotel ratings do not tell the full story of how guests view a hotel, according to an analysis of the text of 5,830 reviews covering 57 hotels in Moscow, Russia. The study found, for instance, that negative comments have a heavier weight in a guest's rating of a hotel than do positive comments. This uneven weighting means that a simple average of positive and negative scores may not provide a clear view of guests' opinion of the hotel. This finding also underlines the importance of consistency, because guests' bad feelings from poor service generally will submerge their favorable feelings from good service. The study applied a regression analysis to the relationships of 18,106 distinct terms relating to five specific attributes—amenities, experience, location, transactions, and value. Reviews for different hotel tiers gave varying weights to those attributes. For instance, the guest's experience was mentioned more commonly in reviews of high-tier hotels, while amenities and location came up more frequently for motels in the middle tier compared to hotels in other tiers. Guests at the lower tier properties wrote more commonly about transactions and value than those staying at hotels in the middle and high tiers. One particularly noticeable feature of the reviews is that ratings sank when guests wrote lengthy reviews that focused tightly on a limited number of hotel attributes, while relatively briefer reviews that took a wider view of the hotel generally had higher ratings. The study also found that guests write more about value and transactions when they are dissatisfied. Thus, text analytics can point to specific steps managers can take to improve guests' assessments and their hotels' ratings.

Keywords

online reviews, data mining, guest reviews, hotel ratings, guest feedback

Disciplines

Hospitality Administration and Management

Comments

What Guests Really Think of Your Hotel: *Text Analytics of Online Customer Reviews*

by Hyun Jeong "Spring" Han, Shawn Mankad, Nagesh Gavirneni, and Rohit Verma

EXECUTIVE SUMMARY

Hotel ratings do not tell the full story of how guests view a hotel, according to an analysis of the text of 5,830 reviews covering 57 hotels in Moscow, Russia. The study found, for instance, that negative comments have a heavier weight in a guest's rating of a hotel than do positive comments. This uneven weighting means that a simple average of positive and negative scores may not provide a clear view of guests' opinion of the hotel. This finding also underlines the importance of consistency, because guests' bad feelings from poor service generally will submerge their favorable feelings from good service. The study applied a regression analysis to the relationships of 18,106 distinct terms relating to five specific attributes—amenities, experience, location, transactions, and value. Reviews for different hotel tiers gave varying weights to those attributes. For instance, the guest's experience was mentioned more commonly in reviews of high-tier hotels, while amenities and location came up more frequently for motels in the middle tier compared to hotels in other tiers. Guests at the lower tier properties wrote more commonly about transactions and value than those staying at hotels in the middle and high tiers. One particularly noticeable feature of the reviews is that ratings sank when guests wrote lengthy reviews that focused tightly on a limited number of hotel attributes, while relatively briefer reviews that took a wider view of the hotel generally had higher ratings. The study also found that guests write more about value and transactions when they are dissatisfied. Thus, text analytics can point to specific steps managers can take to improve guests' assessments and their hotels' ratings.

ABOUT THE AUTHORS

Hyun Jeong "Spring" Han, Ph.D., is an assistant professor in the faculty of business and management at the School of Business Administration at the National University Higher School of Economics in Moscow. Previously a research associate at the Cornell School of Hotel Administration, her research spans international locations including the U.S. and Korea, with studies of tradeshows, conventions, sport-related travel, and health-related hospitality. A graduate of Sungkyungkwan University, she holds a Ph.D. from Kyonggi University.

Shawn Mankad, Ph.D., is an assistant professor of operations, technology and information management at the Samuel Curtis Johnson Graduate School of Management at Cornell University. His research focuses on using data analytics to facilitate better decision-making, through data mining, machine learning, and visualization for economic modeling with unstructured and complex structured data. Previously at the Robert H. Smith School of Business at the University of Maryland, Mankad has been a consultant with the U.S. Commodity Futures Trading Commission and also analyzed market activity for the Federal Reserve Board. A graduate of Carnegie Mellon University, he received his PhD in statistics from the University of Michigan.

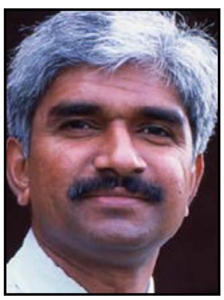

Nagesh Gavirneni, Ph.D., is an associate professor at the Samuel Curtis Johnson Graduate School of Management at Cornell University. His research interests are in the areas of supply chain management, inventory control, production scheduling, simulation, and optimization. His papers have appeared in numerous technical and academic journals. Previously at the Kelley School of Business at Indiana University, he also has held positions at SmartOps, Maxager Technology, and Schlumberger. A graduate of Indian Institute of Technology, he earned his Ph.D. in industrial administration from Carnegie Mellon University and holds master's degrees from Carnegie Mellon University and Iowa State University.

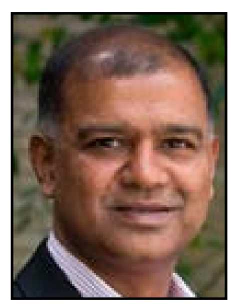

Rohit Verma, Ph.D., is the executive director of the Cornell Institute for Healthy Futures and the Singapore Tourism Board Distinguished Professor in Asian Hospitality Management at the School of Hotel Administration (SHA). He also serves as chair or co-chair for the Cornell Hospitality Research Summit, an industry-academic collaborative conference organized by SHA. Earlier he served as coordinator of the MBA/MMH dual-degree program, director of the Executive Master Program Development Project, and executive director for the Center for Hospitality Research at SHA. Prior to his appointment at Cornell University, Verma was the George Eccles Professor of Management at the David Eccles School of Business, University of Utah. He has taught undergraduate, MBA, and executive courses at several universities around the world including DePaul University, German Graduate School of Business and Law, Helsinki School of Economics, Indian School of Business, Korea University, and the University of Sydney.

tripadvisor®

The authors express their appreciation to TripAdvisor for providing the data for this study.

What Guests Really Think of Your Hotel:

Text Analytics of Online Customer Reviews

by Hyun Jeong "Spring" Han, Shawn Mankad, Nagesh Gavirneni, and Rohit Verma

Customer feedback is an essential source of information for improving operations in the service industry, but capturing an accurate and complete picture of the customer experience has always been a challenging task. With the recent rise of social networks and reviews posted on travel portals, the hotel industry has access to a large and rapidly growing number of online customer reviews, but most of those reviews consist of unstructured comments that are not amenable to direct analysis using traditional methods. In this paper, we demonstrate the use of automated software tools that have been designed to analyze large volumes of text, so that the hotel industry can apply information from these online data to develop actionable managerial insights and improve their operations.

Recognizing the opportunity created by online reviews, a growing number of companies have developed management dashboards that display summaries of online customer reviews, comments on social media, and other information. To learn about these techniques, we interviewed technology officers and visited the websites of companies that were represented at the 2014 Cornell Hospitality Research Summit.[1] We found that many of them use or are actively developing automated approaches to characterize the sentiment found in online reviews. For example, one executive expected to have an automated sentiment function for customers in spring 2016. Vendors are providing such sentiment-assessment tools, including Fishbowl Analytics, which specializes in helping restaurants use data from a variety of sources and recently launched an analytics platform that summarizes sentiment and engagement on social media.[2] Other analytics firms are shown in the appendix at the end of this report. In this research, we apply a suite of natural language processing algorithms that can automatically characterize such dimensions as the style of writing and content from online review and social media channels.

The two main difficulties of extracting information from the text of consumer reviews are the unstructured nature of the data, which impedes direct analysis using standard statistical or econometric techniques, and the (typically) large quantity of data, which renders manual treatment of the data unsuitable. Even if the volume of data could be reasonably coded or interpreted by human agents (e.g., through large-scale crowdsourcing arrangements), a wealth of studies suggests that people are prone to multiple behavioral biases and are unreliable intermediaries if the goal is to draw managerial insights from consumer reviews using a clinical, unbiased, and scientific approach.[3]

To achieve our research objectives and to illustrate our methodologies, we partnered with TripAdvisor, a leading hotel review website, to obtain and analyze posted reviews on hotels in Moscow, Russia, an international hub that we consider to be representative of other global cities. We converted the unstructured text of these reviews to a structured form amenable to analysis, and applied natural language processing algorithms. From our analysis of this dataset, we make the following contributions:

1. We showcase how modern text-mining methods can be applied to hotel reviews to extract valuable information embedded in these reviews. Moreover, the reviews' content can vary substantially (in sentiment, quality of writing, and themes) from the numerical satisfaction ratings assigned by the review writers. This suggests that information from the text can potentially yield insights not indicated in the ratings for how hotels can improve their operations and better meet customer expectations.

2. By measuring each review's positive and negative sentiment, we find that the overall sentiment of a review (positive minus negative sentiment) and the consistency with which the sentiment is expressed in a review are both important factors that influence a review's numerical rating score. Given that it is now possible to easily partition the positive and negative comments in a review, hotel managers should be appropriately prepared to make operational and strategic changes in response to both positive and negative content.

3. Because each review may refer to several different aspects of a hotel's operations, we use topic modeling, an established technique from natural language processing, to determine what components of a guest's stay are addressed in the review. We find that the mere inclusion (or exclusion) of certain topics in a review is significantly associated with that review's numerical rating score. We also show that on average, reviews with higher numerical ratings tend to be shorter and discuss topics more broadly, whereas reviews with lower ratings tend to be longer and focus on a smaller number of major issues.

After a summary of the literature related to hospitality operations, customer reviews, and text analytics, we provide a brief introduction to our analytical methods. After describing our sample and providing some summary statistics and broad findings, we dig into the regression results to derive managerial insights for improving hotel operations and strategy. We conclude by discussing implications and directions for future research.

Methodological versus Managerial Approaches

Studies of text-mining of online consumer reviews tend to fall into one of two categories, the methodological approach and the managerial approach, which give rise to two bodies of literature that are quite distinct in objective and focus. Our study applies the managerial approach, which we review below. In contrast to our research, studies in the methodological literature typically focus on developing new methods and algorithms that, for example, extract new types of features from the text of

---

[1] www.cvent.com/events/cornell-hospitality-research-summit-chrs-2014/speakers-aca238178fcc4d6f99c395f8c1101196.aspx; viewed 8/11/15.

[2] www.czarmetrics.com/fishbowl-launches-next-generation-guest-analytics-platform/; viewed 10/17/2015.

[3] For example, see: Daniel Kahneman (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5), 1449-1475,

consumer reviews or perform computations more efficiently.[4] Well-established methods are described, for instance, by Feldman and Sanger and by Bird, Klein, and Loper, as well as in recent monographs.[5]

Even though methodological studies often employ data from consumer reviews, their primary purpose is to illustrate the particular methodological innovation that the study introduces. In contrast, managerial studies generally apply established methods to generate insights into specific managerial issues. The managerial literature on text-mining online consumer reviews is relatively nascent, and so we briefly trace its inception and development.

Many studies have addressed the role of online reviews in influencing consumer behavior (and, ultimately, sales and revenue). Few such studies have considered the reviews' content, however, instead investigating the effects of ratings or such descriptive statistics as review volume. The results have been mixed. Some studies find that better numerical rating scores are associated with better performance,[6] whereas other studies have found no such associations.[7]

In part because of these mixed findings, researchers turned toward the content of the reviews themselves, seeking to understand their role in consumer decisions and sales performance. Analyses of this type have involved manually transforming text data into quantitative forms by following structured guidelines to mitigate bias and increase objectivity.[8] For example, Pavlou and Dimoka recruited coders to categorize text comments for sellers on eBay to study whether sellers who had better positive feedback were able to command higher prices for their items,[9] and Black and Kelley used this approach to show that hotel reviews that had strong narrative elements of a good story were viewed as more helpful by consumers.[10] However, setting aside the obvious issue of behavioral biases,[11] it's difficult to apply manual methods to large quantities of data. For these reasons, scholars have used automated text mining analysis, as we do in this study.

We take this analysis a step beyond the common approach of making statistical inferences using regression analyses of such measurements as word counts, sentiment scores, and linguistic styles (such as measures of readability). Typical control variables in existing studies are numerical rating scores and other context-specific variables, while dependent variables have included online conversion rates,[12] evaluations of review helpfulness,[13] and a firm's stock market performance.[14] Other studies have used t-tests to compare how these variables differ between deceptive hotel reviews (i.e., those not written by actual consumers) and truthful ones,[15] as well as whether one can use this analytical method to detect fraudulent reviews.[16] A particularly interesting approach was that of Ghose, Ipeirotis, and Li, who built a model to match hotels to consumer preferences by combining

---

[4] For example, see: Titov, I., & McDonald, R. (2008). Modeling online reviews with multi-grain topic models, in: Proceedings of the 17th international conference on World Wide Web, Beijing, China, 111-120; Baccianella, S., Esuli, A., & Sebastiani, F. (2009). Multi-facet rating of product reviews, in: *Advances in Information Retrieval*, ed. Boughanem, M., Berrut, C., Mothe, J., & Soule-Dupuy, C. Springer, Berlin, 461-472; Wang, H., Lu, Y., & Zhai, C. (2010). Latent aspect rating analysis on review text data: A rating regression approach, in: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, Washington DC, 783-792; and McAuley, J., & Leskovec, J., (2013). Hidden factors and hidden topics: Understanding rating dimensions with review text. Proceedings of the 7th ACM conference on Recommender systems, Hong Kong, 165-172.

[5] Feldman, R., & Sanger, J. (2007). *The text mining handbook: Advanced approaches in analyzing unstructured data*. Cambridge University Press, New York; Bird, S., Klein, E., & Loper, E. (2009). Natural language processing with Python. O'Reilly Media Inc., Sebastopol, CA; *Handbook of latent semantic analysis*. ed. Landauer, T.K., McNamara, D.S., Dennis, S., & Kintsch. Routledge: Taylor and Francis Group, New York, 2013; and Liu, B. (2010). Sentiment Analysis and Subjectivity, in *Handbook of Natural Language Processing*, 2nd Edition. ed. Indurkhya, N., & Damerau, F.J. Chapman & Hall/CRC Machine Learning & Pattern Recognition, 627-666.

[6] For example: Chevalier, J. A., & Mayzlin, D. (2006). The Effect of Word of Mouth on Sales: Online Book Reviews. *Journal of Marketing Research*, 43(3) 345-354; and Dellarocas, C., Awad, N.F., & Zhang, M. (2004) Using Online Reviews as a Proxy of Word-of-Mouth for Motion Picture Revenue Forecasting. Working Paper. Available at ssrn.com/abstract=620821.

[7] For example: Chen, P., Wu, S., & Yoon, J. (2004). The Impact of Online Recommendations and Consumer Feedback on Sales. Proceedings of the International Conference on Information Systems, Washington DC, 711–724; and Blal, I. and Sturman, M.C. (2014). The Differential Effects of the Quality and Quantity of Online Reviews on Hotel Room Sales. *Cornell Hospitality Quarterly*, 55(4), 365-375.

[8] For example: Kassarjian, H.H. (1977). Content Analysis in Consumer Research. *Journal of Consumer Research*. 4(1), 8-18.

[9] Pavlou, P.A., & Dimoka, A. (2006). The Nature and Role of Feedback Text Comments in Online Marketplaces: Implications for Trust Building, Price Premiums, and Seller Differentiation. *Information Systems Research*, 17(4), 391-412.

[10] Black, H. G., & Kelley, S. W. (2009). A storytelling perspective on online customer reviews reporting service failure and recovery. *Journal of Travel & Tourism Marketing*, 26(2), 169-179.

[11] For example: Kahneman, *op.cit*.

[12] Ludwig, S., de Ruyter, K., Friedman, M., Brüggen, E. C., Wetzels, M., and Pfann, G. (2013). More than words: The influence of affective content and linguistic style matches in online reviews on conversion rates. *Journal of Marketing* 77(1), 87-103.

[13] Korfiatis, N., García-Bariocanal, E., & Sánchez-Alonso, S. (2012). Evaluating content quality and helpfulness of online product reviews: The interplay of review helpfulness vs. review content. *Electronic Commerce Research and Applications*, 11(3), 205-217.

[14] Tirunillai, S., & Tellis, G. J. (2012). Does chatter really matter? Dynamics of user-generated content and stock performance. *Marketing Science*, 31(2), 198-215.

[15] Yoo, K. H., & Gretzel, U. (2009). Comparison of deceptive and truthful travel reviews, in: *Information and communication technologies in tourism*, ed. Hopken, W., Gretzel U., & Law, R. Springer-Verlag, The Netherlands, 37-47.

[16] Hu, N., Koh, N.S., Reddy, S.K. (2014). Ratings lead you to the product, reviews help you clinch it? The mediating role of online review sentiments on product sales. *Decision Support Systems* 57, 42-53.

text analysis with other data, including hotel location information and crowdsourced, user-contributed opinions.[17]

We know of two studies that took approaches similar to the one we use here, although we use a more sophisticated analysis. First, Ghose and colleagues extended their base model by clustering the nouns and noun phrases of reviews to generate additional features for their regression analyses.[18] Second, Lee and Bradlow applied K-means clustering to infer product attributes of digital cameras from the pro-and-con lists accompanying their online reviews.[19] Like those works, our study aims to uncover a set of attributes or themes that run through the entire set of reviews. At the same time, our approach is distinct from those clustering methods because it is based on a technique known as Latent Dirichlet Allocation (LDA). As presented by Blei and Lafferty, LDA results in more informative summaries of text data than other approaches. [20] This is primarily because the LDA model is founded on more realistic assumptions and has greater flexibility. For example, LDA can represent each text with multiple topics instead of a single topic.

Fundamental Concepts in Text Analysis

In this section, we briefly review some of the popular analytical techniques associated with text analytics. We do this in order to lay the groundwork for our study, as well as describe our methodology and introduce terms that we use in this paper. The techniques we use involve preprocessing and text representation, sentiment analysis, and writing style evaluation.

Preprocessing and text representation. Arguably the most critical step in analyzing unstructured text documents is transforming free form text into a structured form that is amenable to analysis. The most popular transformation is the so-called "bag of words" representation of text. [21] This involves representing the set of documents to be analyzed (usually referred to as a "corpus") with a document-term matrix,[22] which contains a column for each word that appears anywhere in the corpus and a row for every document (in this case, each consumer review). Each matrix entry is a count of the number of times a particular word appears in each document.

The document-term matrix is a structured table of numbers that can in principle be analyzed using standard techniques. However, in practice the matrix often grows to be too large, creating computational and memory challenges. Thus, analyses based on text mining require preprocessing of the corpus to retain meaningful words and remove uninformative ones, thereby keeping the number of terms that appear in the corpus from becoming excessive. Standard preprocessing steps are: **(1)** transforming all text into lowercase, **(2)** removing words composed of fewer than three characters and so-called stop words (e.g., the, and, of), **(3)** stemming words by removing suffixes, so that words like values, valued, and valuing are all replaced with valu, and, finally **(4)** removing words that occur either too frequently or rarely.[23]

Sentiment analysis. The word count and sentiment represent the most basic statistics for summarizing a corpus, especially since research has shown that they are associated with customer decision making and product sales.[24]

To investigate potential non-linearities in the impact of sentiment we separately computed measures of each review's positive sentiment, $(\text{Sentiment})^+$, and negative sentiment, $(\text{Sentiment})^-$. The $(\text{Sentiment})^+$ measure was calculated by counting the number of words in the review that matched a list of positive words in validated databases (known as dictionaries), and $(\text{Sentiment})^-$ was calculated with negative dictionary words.

The choice of dictionary is an important methodological consideration, since the sentiment relating to certain words changes with the underlying context. So, for instance, dictionaries created using financial 10-K disclosures might not be appropriate for areas other than financial sentiment analysis, while other dictionaries would not work as well as a financial dictionary.[25] Not having a specific dictionary for our topic, we came as close as we could by choosing dictionaries used by Liu, by Nielsen, and by Mohommad and Turney.[26] These were created to summarize the opinions within online customer reviews and to perform tonal analysis of social media blogs. In total, the combined dictionaries consist of approximately 10,000 words.

Writing style. Measures of readability (also known as reading levels) are numerical scores that indicate how easy

[17] Ghose, A., Ipeirotis, P. G., and Li, B. (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science* 31(3), pp. 493-520.

[18] *Ibid.*

[19] Lee, T. Y., & Bradlow, E. T. (2011). Automated marketing research using online customer reviews. *Journal of Marketing Research*, 48(5), 881-894.

[20] Blei D. and Lafferty J. (2009). Topic Models, in *Text Mining: Classification, Clustering, and Applications*, ed. Srivastava, A.N., & Sahami M. CRC Press, Boca Raton, FL, 71-94.

[21] The moniker "bag of words" stems from the assumption that the distribution of words within each document is sufficient. That is, linguistic features like word order, grammar, and other attributes within written text can be safely ignored for statistical analysis.

[22] Kosala, R., & Blockeel, H. (2000). Web mining research: A survey. , 2(1), 1-15.

[23] Boyd-Graber, J., Mimno, D., & Newman D. (2014). Care and Feeding of Topic Models: Problems, Diagnostics, and Improvements, in *Handbook of Mixed Membership Models and Their Applications*, ed. Airoldi E.M., Blei, D.M., Erosheva E.A., & Feinberg, S.E. CRC Press, Boca Raton.

[24] Hu *et al. op.cit.*

[25] Loughran, T. and McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35-65.

[26] Liu, op.cit.; Nielsen, F.Å. (2011). A new ANEW: Evaluation of a word list for sentiment analysis in microblogs, in Proceedings of the ESCW2011 Workshop on 'Making Sense of Microposts': Big things come in small packages, ed. Rowe M., Stankovic M., Dadzi A, & Hardey M. Crete. 93-98.; and Mohammad, S. M. and Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3), 436-465.

or difficult a given piece of text is to comprehend. To measure readability we employed the well established Flesch-Kincaid Grade Level tool.[27] To estimate the comprehension difficulty of a given piece of text, this measure combines weighted ratios of word count to number of sentences and number of syllables to total words. Available within Microsoft Word, the Flesch-Kincaid tool is popular due to its fast calculation and history of reliable results.[28]

Topic Modeling using LDA

Topic modeling algorithms automatically summarize large archives of text by discovering hidden topics or themes found within a set of documents.[29] As we mentioned above, Latent Dirichlet Allocation is a powerful and widely used topic modeling algorithm in which the hidden topic structure is inferred from the original texts using a probabilistic framework. The idea behind this method is that all documents share the same topic set, but each document exhibits a different probabilistic mixture of those topics. In short, certain words are more likely to be used with certain topics. LDA employs a Bayesian estimation framework to the text to infer the topics (distributions of words) and decomposes each document into a mixture of topics. Specializing to our setting, the outputs of LDA are two probability distributions: *P(topic|review)*, which is the probability distribution of topics for a given review, and *P(word|topic)*, which is the probability distribution of words for a given topic.

From the probability distributions that were output by the LDA procedure, we computed two sets of variables for each review *i*. First, for each topic, we created a dummy variable, $Topic_{ik}$, which was equal to 1 if the posterior probability $P(k|i)$ was greater than a threshold, and 0 otherwise. The threshold we used was the median probability across all reviews and topics. Thus, we interpret $Topic_{ik} = 1$ to mean that review *i* discusses topic *k*. Second, we created a continuous variable called focus,[30] which we define as $Focus_i = \sum_{(k=1)}^{K} P(k|i)^2$. Reviews with larger values of focus are those that concentrate on a relatively small number of topics.

Data and Methods

We obtained all customer reviews on TripAdvisor for hotels in Moscow, Russia, from January 1, 2012, through December 31, 2013. For each review, we observed the date, customer satisfaction rating scores (1–5), title and full text of the review, and type of traveler (i.e., business, family, solo, couple). Because we were prohibited from identifying the hotels in our dataset (in particular, we could not link the hotels in our dataset to their star rating), we used the guests' numerical ratings as our primary measure of service quality or excellence in all our analyses. We analyzed the reviews that **(1)** were written in English, **(2)** were written for hotels that had 30 or more reviews during this period, and **(3)** were not duplicates, for a total of 5,830 reviews covering 57 hotels (see Exhibit 1).

Exhibit 1

Sample construction

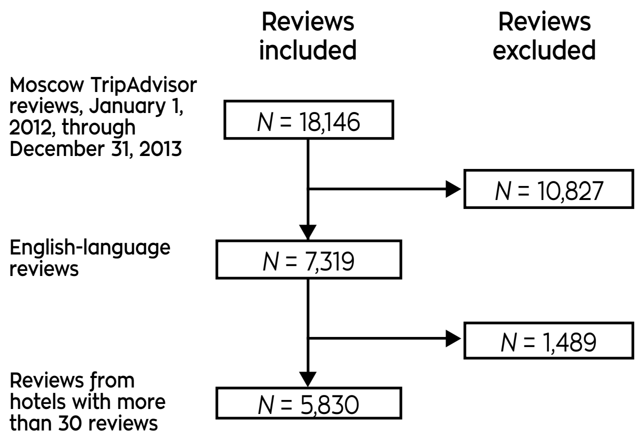

We first analyzed the reviews for individual hotels, segmented into three equal tiers (low, middle, and high), according to the tertiles of the mean rating from all reviews for that hotel. We computed summary statistics (means and standard errors) for various text-based quantities, as described below. The objective of this level of analysis was to better understand the structure of the data and to identify broad trends. We then analyzed individual reviews using linear regression analyses to study the relationship of various text-based review features on rating scores, after controlling for other variables.

Preprocessing and text representation. We employed the standard preprocessing steps described above: converting to lowercase, removing stop words, stemming words, and removing words that occurred frequently (namely, hotel, room, good, Moscow, and stay) or rarely (in this case, ten times or less). This step resulted in a document-term matrix containing 5,830 documents (i.e., the number of reviews in our sample), and 18,106 distinct terms.

---

[27] Kincaid, J.P., Fishburne, R.P., Rogers R.L., and Chissom, B.S. (1975). Derivation of new readability formulas (Automated Readability Index, Fog count, and Flesch reading ease formula) for Navy enlisted personnel. Research Branch Report 8-75, Chief of Naval Technical Training: Naval Air Station, Memphis.

[28] *Ibid.*

[29] Blei and Lafferty, *op. cit.*

[30] Also known as the Herfindahl Index. See: Rhoades, S.A. (1993). The Herfindahl-Hirschman index, *Federal Reserve Bulletin*, March 1993 pp. 188-189.

Exhibit 2

Examples of reviews with lowest and highest readability scores

Grade level	Review Text
Lowest possible, less than 10	"Best location, well trained friendly staff. Nice restaurant on the top. Clean rooms. What else do you need :) ?"
	"Short and sweet….if you can afford to stay here, stay here"
	"i was in hilton hotel in july and i was so nice trip. quality of rest was high. if you want to have fun you should go to hilton."
	"We stayed 4 days with friends. As we hat late flight we came in Da Hostel deep night. The check-in was fast. Rooms are not so big but clean and nice. And we were so lucky to have friendly neighbors from France. We spent nice holidays all together. And we liked design so much!"
Highest possible, greater than 16	"Beautiful historical hotel conveniently located. The staff are very friendly and accommodating. The price is very reasonable. We've enjoyed our stay very much! Highly recommended to everybody traveling to Moscow."
	"A typical Swedish style hotel on the banks of the Volga—very bright, clean and airy—with a free minibar! Continental style breakfast excellent A fair way out from the city centre but not too bad once you have mastered the metro There is a very good pizza restaurant with reasonably priced wine at the end of the road"
	"Excellent historic hotel in a great location, near several train stations and metro stops. Beautiful lobby, comfortable rooms. Certainly a traveler expecting Western standards in a business hotel will not leave here disappointed."
	"Expensive but very luxurious, perfect retreat after a full day of meetings. Service and facilities surpassed all expectations. Not far from many restaurants, try the Azerbaijani restaurant around the corner from the hotel."

Sentiment analysis. Our sentiment analysis found that reviews tended to have less negative content than positive content. The median negative sentiment was -3 with a maximum (magnitude) of -47, whereas the median positive sentiment was +9 with a maximum of +147. The mean negative sentiment was -3.02 with a standard deviation of 3.65, whereas the mean positive content was 10.65 with a standard deviation of 7.22.

Writing style. Our assessment of the reviews' grade level estimated that the writing style of around 71.1 percent of the reviews was below high school level, while just 6.3 percent of the reviews had a quality indicative of a college degree or higher. Exhibit 2 shows, for illustrative purposes, reviews that had the lowest and highest estimated grade levels.

Topic modeling. Our topic modeling used the software implementation of LDA and guidelines on parameter values by Gruen and Hornick,[31] and we applied Gibbs sampling for the main computation. Because we had to determine the number of

[31] Gruen B. and Hornik K. (2011). Topic Models: An R Package for Fitting Topic Models. *Journal of Statistical Software* 40(13), 1-30. www.jstatsoft.org/v40/i13/.

Exhibit 3

Cross-validated perplexity score against number of topics

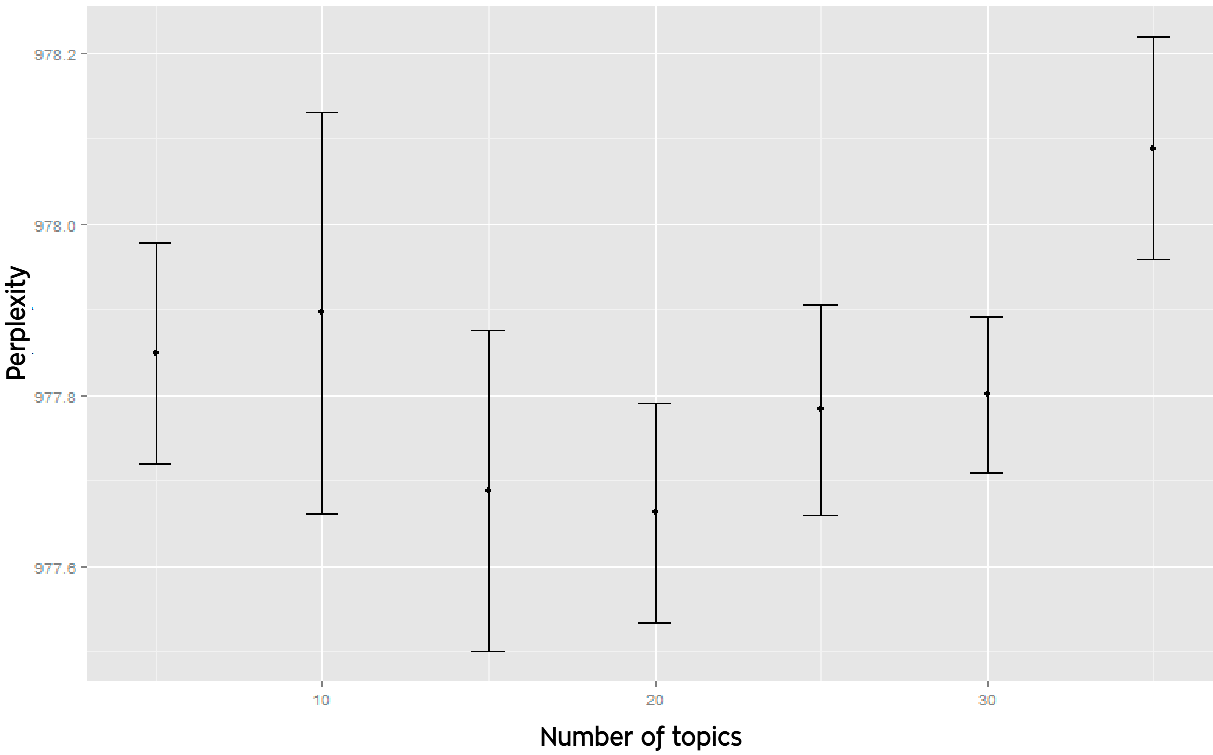

Note: Error bars represent 95% confidence intervals around the mean.

topics in advance, we assessed the perplexity score (a goodness-of-fit measure) to choose this number. A lower perplexity value represents a better fit. When we evaluated perplexity scores using 5-fold (randomized) cross-validation with ten repetitions for 5, 10, 15, 20, 25, 30, and 35 topics, we found little variance in the perplexity scores (as shown in Exhibit 3). Therefore, we decided to use the smallest number of topics—five—for simplicity and increased interpretability. The topics are amenities, location, transactions, value, and experience.

Exhibit 4

Most likely words in each topic

Term	*Topic 1* Amenities	*Topic 2* Location	*Topic 3* Transactions	*Topic 4* Value	*Topic 5* Experience
1	breakfast	metro	check	one	great
2	free	walk	night	like	staff
3	food	locat	time	price	servic
4	bed	station	day	can	excel
5	floor	clean	back	time	locat
6	bar	minut	even	place	nice
7	wifi	close	book	will	red
8	small	also	ask	much	best
9	bathroom	citi	front	just	view
10	night	restaur	desk	better	recommend
11	well	nice	one	busi	location
12	also	staff	servic	standard	well
13	restaur	airport	first	bit	help
14	area	street	arriv	star	perfect
15	larg	english	recept	get	visit
16	water	train	got	mani	high
17	buffet	can	get	realli	veri
18	clean	near	taxi	need	kremlin
19	shower	away	just	quit	realli
20	use	just	made	expect	square
21	includ	center	will	internet	breakfast
22	nice	busi	hour	lobbi	love
23	offer	shop	make	want	friend
24	etc	get	call	old	trip
25	day	friend	never	look	enjoy
26	english	around	russian	thing	spacious
27	comfort	min	guest	feel	service
28	drink	speak	next	littl	definit
29	smoke	conveni	way	big	beauti
30	coffe	red	took	lot	place

The LDA topic modeling algorithm outputs posterior probabilities that capture how likely it is that a given word belongs to a particular topic. Using these probabilities we ranked the thirty most likely words in each topic in order of how likely it was that each word belongs to that topic (as shown in Exhibit 4). Typical words for the five topics are as follows:

- *Topic 1*: ***Amenities***. The likely terms associated with the amenities provided by the hotel include breakfast, free, food, bed, and wi-fi.
- *Topic 2*: ***Location***. Terms associated with the location or vicinity of the hotel include metro, walk, locat (a stemmed word that includes location and located), station, and restaurant.
- *Topic 3*: ***Transactions***. Transaction terms connected to the mechanics of a guest's stay include one, check, night, time, day, arriv (including arrival and arriving), front, and desk.
- *Topic 4*: ***Value***. Terms associated with guests' perceived value or money include price, like, get, better, look, find, standard, and internet.
- *Topic 5*: ***Experience***. Finally, terms associated with a guest's overall experience include generally positive descriptive terms, such as great, excel(lent), and recommend, and terms that pertain to particular characteristics, such as staff, view, high, and pool.

To make these terms more useful for management action, one could conceptualize these topics by putting them on a timeline and interpreting them according to when these topics might be salient during a guest's interactions with the hotel. For simplicity, one could conceptualize five distinct periods: **(1)** before arrival, **(2)** check-in, **(3)** stay, **(4)** check-out, and **(5)** after departure. If a particular hotel receives negative reviews that focus on a certain topic, then a manager can use this classification to assess guests' interactions with the hotel and determine operational deficiencies.

As examples of how the five periods connect with the topic analysis, we relate the topics to the periods (as shown in Exhibit 5). Location (Topic 2) is likely to be important to a guest before arrival and after departure as she travels to and from the hotel, and the locality is also important for her during her stay. Transactions (Topic 3) are likely most salient during check-in and check-out, and to a certain extent, even before she arrives at the hotel, when she is booking her stay. Value (Topic 4) is likely to be important to her before arrival when she is choosing which hotel to book, and even after the stay when she reflects on the totality of the experience and considers whether this stay was a good value for her money.

Another way to conceptualize these topics is to categorize them by the types of choice they represent—whether strategic or operational. Two topics seem clearly related to strategic choices, namely, location, which has to do with an upfront strategic choice about where the hotel is situated, and value, which is related to the hotel's price position. Two other topics are clearly related to operational choices: experience relates to the operational decisions taken by managers to engage with the guest's wants and needs during her stay, and transactions relates to operational decisions that can either create or reduce frictions in the mechanics of her stay. Finally, amenities seem to be related to both strategic and operational choices, since amenities relate to a hotel's market positioning and competitive strategy, but also depend on how well amenities are maintained and presented each day. Given these connections, reviews can represent a source of valuable information for managers to assess the impact of the choices that they make, and this categorization of topics can serve as a way to structure this information so that managers have a better understanding of the nature of this assessment.

Regression models. To analyze the reviews, our regression models evaluated the associations between the numerical rating, hotel tier, customer type, negative and positive components of the sentiment score, style of writing, and topic inclusion. The dependent variable was the numerical rating score of the review for each model, with the other terms as independent variables (see Exhibit 6). To aid interpretation, we used the logarithm of word count to correct a skewed distribution. We further standardized all continuous variables except for positive and negative sentiment. We included hotel fixed effects in all models and calculated heteroscedasticity-adjusted robust standard errors clustered at the hotel level,[32] because we expected significant variation between hotels. To assess multicollinearity in our models, we calculated the (generalized) variance inflation factors of each regressor variable in our model.[33] In addition, to assess the sensitivity of our results to those of the Flesch-Kincaid index of readability, we also re-ran the models using several other common readability indices.[34]

Exhibit 5

Illustration of five periods of timeline and relevance of topics in each period

	Before Arrival	Check-in	Stay	Check-out	After Departure
Topic 1: Amenities			✓		
Topic 2: Location	✓		✓		✓
Topic 3: Transactions	✓	✓		✓	
Topic 4: Value	✓				✓
Topic 5: Experience		✓	✓	✓	

Exhibit 6

Description of regression models

Model	Independent Variables
1	Hotel level fixed effects, Log(Word count), $(\text{Sentiment})^+$, $(\text{Sentiment})^-$
2	Variables from Model 1, $[(\text{Sentiment})^+]^2$, $[(\text{Sentiment})^{-2}$
3	Variables from Model 1, $[(\text{Sentiment})^+]^2$ + $[(\text{Sentiment})^-]^2$
4	Variables from Model 3, focus
5	Variables from Model 4, Flesch-Kincaid Grade Levels
6	Variables from Model 5, dummy variables for traveler types
7	Variables from Model 6, dummy variables for topics
8	Variables from Model 7, interaction terms between topic dummies and $(\text{Sentiment})^+$, interaction terms between topic dummies and $(\text{Sentiment})^-$

[32] White, S. (2003) "The 2003 National Assessment of Adult Literacy (NAAL)," Center for Education Statistics (NCES), Technical Report NCES 2003495rev, US Dept. of Education, Inst. of Education Sciences, nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2003495rev.

[33] Fox, J., & Monette, G. (1992). Generalized collinearity diagnostics. *Journal of the American Statistical Association* 87(417), pp. 178-183.

[34] See: DuBay, W.H. *The Principles of Readability*. Impact Information, www.nald.ca/library/research/readab/readab.pdf, 2004.

Exhibit 7

Number of reviews by each traveler type as a percentage of the total number of reviews in each hotel tier

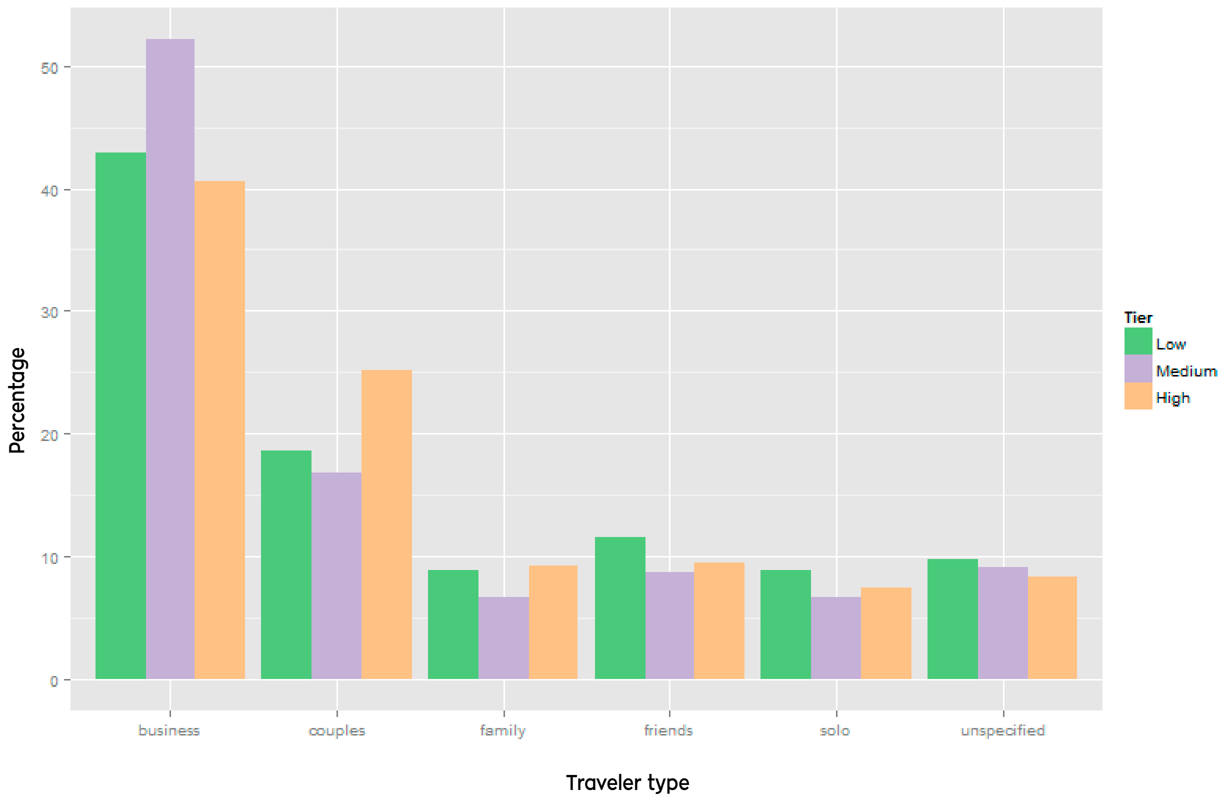

Discussion of Summary Statistics

The hotels' average rating scores were 4.4. for the high tier, 4.0 for the middle tier, and 3.5 for the low tier. We found that hotels in the high rating tier had 147 reviews on average, compared to those in the low tier, which averaged 65 reviews.

Business travelers constituted the largest proportion of travelers who wrote reviews, exceeding 40 percent in all three tiers. The middle tier had the largest percentage of business travelers, comprising slightly more than 52 percent of the reviews. As illustrated in Exhibit 7, couples had the second largest percentages across all the three tiers (25% for the upper tier, 17% for the middle tier, and 19% for the lower tier). Approximately 10 percent of the reviews had no information on the customer type.

Exhibit 8

Number of reviews mentioning each topic, expressed as a percentage of total number of reviews in each hotel tier

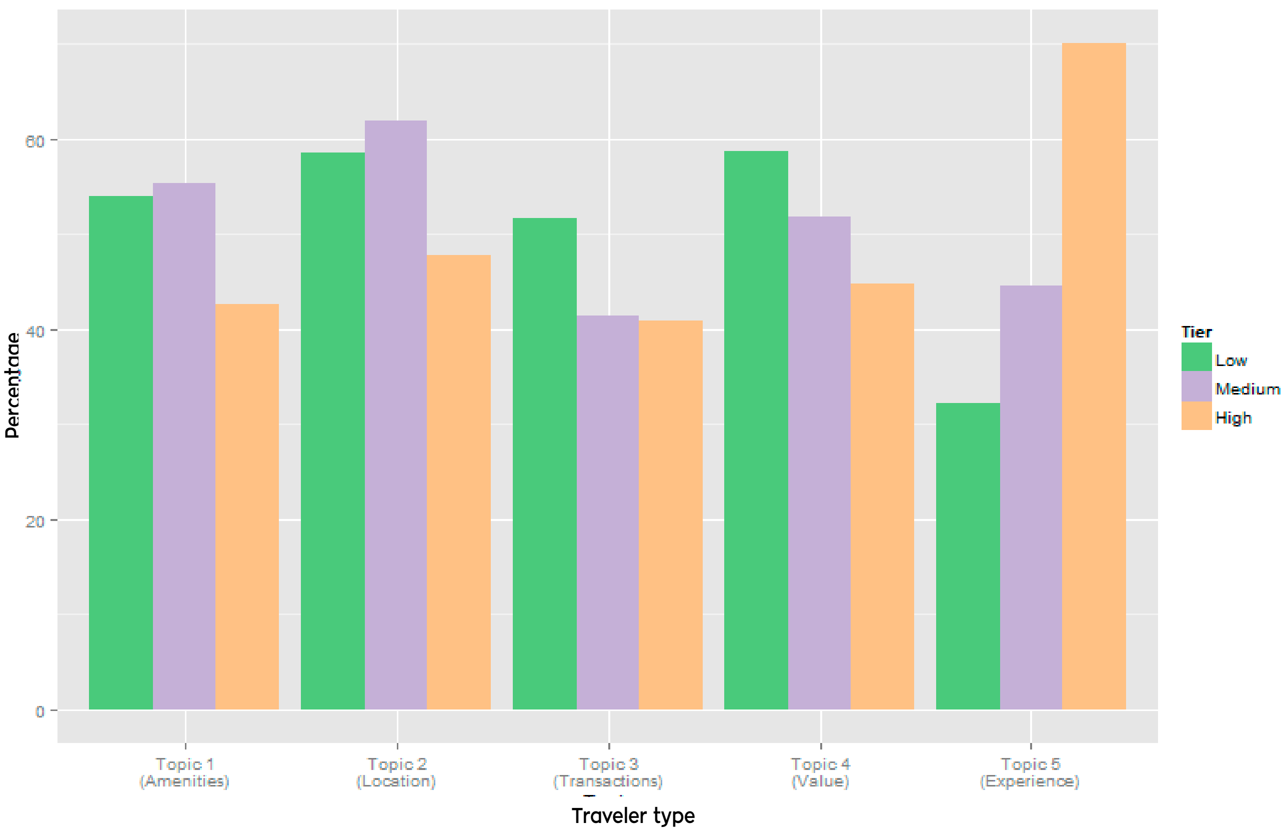

We observed several interesting differences in the reviews for the various hotel tiers. Experience (Topic 5) strongly dominated reviews for high-tier hotels, with approximately 70 percent discussing the guest's experience. In contrast, experience was mentioned in only 32 percent of the reviews in the low tier and 45 percent of reviews in the middle tier. Amenities (Topic 1) and location (Topic 2) came up more frequently for motels in the middle tier compared to hotels in the low and high tiers. In contrast, hotels in the low tier had larger proportions of reviews written about transactions (Topic 3) and value (Topic 4) compared to hotels in the middle and high tiers.

Exhibit 9a

Regression results (part 1)

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Intercept	3.240 (0.044)***	3.147 (0.041)***	3.041 (0.049)***	3.049 (0.047)***	3.048 (0.047)***	2.981 (0.051)***	3.136 (0.062)***	3.109 (0.076)***
Log(word count)	-0.203 (0.026)***	-0.223 (0.021)***	-0.270 (0.024)***	-0.231 (0.022)***	-0.232 (0.022)***	-0.239 (0.022)***	-0.116 (0.020)***	-0.086 (0.020)***
$(\text{Sentiment})^-$	-0.113 (0.006)***	-0.150 (0.006)***	-0.098 (0.006)***	-0.093 (0.006)***	-0.093 (0.006)***	-0.091 (0.006)***	-0.076 (0.006)***	-0.091 (0.012)***
$(\text{Sentiment})^+$	0.067 (0.004)***	0.091 (0.005)***	0.088 (0.005)***	0.086 (0.004)***	0.086 (0.004)***	0.085 (0.004)***	0.069 (0.004)***	0.080 (0.007)***
$[(\text{Sentiment})^-]^2$		0.003 (0.000)***						
$[(\text{Sentiment})^+]^2$		-0.001 (0.000)***						
$[(\text{Sentiment})^-]^2 + [(\text{Sentiment})^+]^2$			-0.000 (0.000)***	-0.000 (0.000)***	-0.000 (0.000)***	-0.000 (0.000)***	-0.000 (0.000)***	-0.000 (0.000)***
Focus				-0.092 (0.013)***	-0.092 (0.013)***	-0.093 (0.013)***	-0.104 (0.013)***	-0.090 (0.015)***
Flesch-Kincaid Grade Level					0.009 (0.011)	0.008 (0.011)	0.009 (0.011)	0.010 (0.011)

Notes: Mean and standard errors shown in (parentheses), segmented by hotel tier. *** = *p*-value < 0.001, ** = *p*-value < 0.01, * = *p*-value < 0.05.

Regression Results

Sentiment. Not surprisingly, reviews with stronger negative sentiment scores were associated with lower ratings in all models, and those with more positive sentiment had higher ratings (see Exhibits 9a and 9b). But the matter is not as simple as that. One important observation is that we found that negative sentiment had a greater impact on scores than did positive sentiment (Exhibit 9a, Model 1). This is consistent with other studies of consumer reviews that have reported asymmetric effects of positive and negative sentiment on various metrics.[35] In this study, all else being equal, we found that an additional negative word in a review was associated with about a 0.11-point decline in rating score, whereas an additional positive word was associated with about a 0.09-point increase.

When quadratic nonlinearities of sentiment were separately entered as regressors in Model 2, we found that their estimated coefficients were reversed in sign from the linear terms, suggesting that the sensitivity of rating scores to sentiment was less pronounced at more extreme values of sentiment.[36]

[35] For example: Yoo and Gretzel, *op.cit.*; and Ludwig *et al., op.cit.*

[36] Consistent with: Ludwig *et al., op.cit.*

Our results highlight two key issues regarding customer reviews. First, the asymmetric impact of positive and negative sentiment on ratings suggests that a simplistic approach that only calculates an "overall" sentiment by taking the difference of positive and negative sentiment does not adequately capture the nonlinear impact of these components. Second, we see the effects of volatility in the reviews, as shown in Model 1. Say that we compare two reviews that both have the same overall sentiment, but review A has a larger variation in sentiment (i.e., contains more positive sentiment and also more negative sentiment than B). In that situation, review A will tend to have a lower rating score than review B. This idea is supported by the findings of Models 3 through 8, where the coefficients of positive and negative sentiment were much more similar when the sum of squares of both positive and negative sentiment were added as regressors. Therefore, our findings consistently suggest that the overall sentiment is a distinct construct from the variation in sentiment associated with ratings.

These findings argue for operational consistency. It is better for hotels to provide guests with a moderately good overall experience than an experience that is extremely good in some regards and terrible in other ways, because in terms of ratings

Exhibit 9b

Regression results (part 2)

		Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Traveler Type:	Couples						0.153 (0.029)***	0.117 (0.026)***	0.114 (0.025)***
	Family						0.190 (0.046)***	0.139 (0.043)**	0.135 (0.041)***
	Friends						0.096 (0.038)*	0.063 (0.036)	0.067 0.036)
	Solo						0.118 (0.043)**	0.106 (0.039)**	0.091 (0.038)*
	Unspecified						0.060 (0.034)	0.054 (0.036)	0.049 (0.036)
Topics	Topic 1 (amenities)							-0.067 (0.021)**	-0.006 (0.044)
	Topic 2 (location)							0.121 (0.021)***	0.182 (0.035)***
	Topic 3 (transactions)							-0.122 (0.026)***	-0.253 (0.037)***
	Topic 4 (value)							-0.177 (0.025)***	-0.268 (0.040)***
	Topic 5 (experience)							0.307 (0.025)***	0.571 (0.046)***
(Sentiment)$^{-}$ x Topics	(Sentiment)$^{-}$ x Topic 1								0.027 (0.006)***
	(Sentiment)$^{-}$ x Topic 2								0.030 (0.011)**
	(Sentiment)$^{-}$ x Topic 3								-0.013 (0.007)*
	(Sentiment)$^{-}$ x Topic 4								-0.009 (0.009)
	(Sentiment)$^{-}$ x Topic 5								-0.008 (0.009)
(Sentiment)$^{+}$ x Topics	(Sentiment)$^{+}$ x Topic 1								-0.013 (0.004)**
	(Sentiment)$^{+}$ x Topic 2								-0.016 (0.004)***
	(Sentiment)$^{+}$ $^{+}$x Topic 3								0.016 (0.004)***
	(Sentiment)$^{+}$ x Topic 4								0.010 (0.003)**
	(Sentiment)+ x Topic 5								-0.024 (0.004)***
R^2		0.435	0.454	0.442	0.45	0.45	0.454	0.489	0.508
Adjusted R^2		0.429	0.448	0.437	0.444	0.444	0.448	0.483	0.501

Notes: Mean and standard errors shown in (parentheses), segmented by hotel tier. *** = *p*-value < 0.001, ** = *p*-value < 0.01, * = *p*-value < 0.05.

the weight of the terrible service will swamp the good feelings from the stay's excellent aspects. Therefore, in addition to the overall quality of service, the consistency with which this service is delivered to a guest is itself an important driver of guest satisfaction.

Word count and focus. We observed that even though unhappy guests were more loquacious, they nevertheless zoomed in on the source of their concern. Across all models, we found a negative association between ratings and word count, as well as a negative association between rating scores and focus. Longer reviews tended to accompany poorer ratings. Likewise, if the review focused on just a few topics, it was more likely to carry a low numerical rating.

Topics. Finally, we noted associations between review topics and overall ratings. Reviews that were written about location and experience were significantly associated with higher ratings in models 7 and 8, where topics were added as regressors (Exhibit 9B). On the other hand, reviews that focused on transactions or value were significantly associated with lower ratings in these models. A discussion of amenities was associated with significantly lower ratings, but only in Model 7. These results suggest that on average, guests tend to write more about "value" and "transactions" when they are dissatisfied, but tend to write more about "location" and "experience" when they are satisfied.

By delving deeper into the connections between topics and sentiment in Model 8, we observe interesting interactions between topics and sentiment. For amenities and location, the coefficients of their interaction terms (with negative and positive sentiment) were reversed in sign from their linear coefficients. This means that the relationship between sentiment and ratings becomes less pronounced in reviews that discuss amenities and location. We might think of these as "objectivity-enhancing" topics, because reviews that mention these topics have ratings that are more independent from the sentiment of their review. Conversely, for reviews that zero in on the transactions topic, the coefficient of its interaction term with negative and positive sentiment had the same sign as the corresponding linear sentiment coefficients. This is the opposite effect from amenities and location, in that reviews that mention those topics tend to have rating scores that have an accentuated effect of sentiment. Hence, one might think of transactions as being an "objectivity reducing" topic. This type of effect is not seen in reviews discussing value or experience, where only the coefficients of their interaction terms with positive sentiment were significantly different from zero.

Conclusion and Implications

Despite this study's limitations, we see implications from this study for customers, managers, and the travel portals themselves. Even though customer feedback is vital for the hotel industry's efforts toward continuous improvement, a comprehensive characterization of the customer experience remains difficult to achieve, in spite of (or perhaps because of) the immense volume of guest commentary on the web. The ability to analyze text contained in online reviews presents an opportunity for greater clarity, as we demonstrate in this study. We show that text-mining algorithms that **(1)** quantify the review sentiment, **(2)** identify the emotional content, and **(3)** extract the main topics of discussion can provide valuable support to management.

Hotel customers. Customers should understand the effect of review sentiment and focus on the ratings found on travel portals, social networks, and the customer reviews. At minimum, we find that the ratings alone do not always capture the guests' actual experiences, and would-be guests should read the reviews with that in mind. Our study also points to ways to augment the value of reviews for selecting a hotel and for encouraging hotel managers to make improvements proposed by reviews. For the customers to ensure the power of their reviews, they must **(1)** participate in large numbers; **(2)** write reviews with high expositional quality; **(3)** ensure that the topics are easily evident; and **(4)** be explicit about their sentiment.

Facility operators. Along that line reviews can be a valuable resource as hotel operators seek to improve their operational effectiveness. We see online reviews as being complementary to the insights obtained from traditional guest satisfaction surveys. Because these data can be voluminous and unstructured, we suggest additional text-mining methods, including the approach outlined here. Among other implications, our results indicate that hotel operators should **(1)** pay special attention to the customer type; **(2)** note what each type of customer seeks in a hotel stay; **(3)** elicit reviews from a larger portion of customers; **(4)** act on the negative feedback to improve their offerings; and **(5)** take advantage of positive feedback to better market themselves.

Travel portals and social networks. Travel websites already ask customers to submit reviews, and we encourage the portals to ensure that the reviews are in a format that is amenable to analysis. They should **(1)** encourage more customers to submit reviews; **(2)** obtain more information from customers (e.g., travel purpose); and **(3)** provide tools to ensure superior quality of reviews.

Limitations and future research. We see several directions for future research, some of them based on our study's limitations. First, our analysis relies only on reviews from TripAdvisor. While we are confident that we had a large enough dataset to support our conclusions, it is conceivable that reviews on that site have some type of invisible bias. To correct for that possibility, future studies should analyze reviews from other online review websites and social media sites. Second, our data were drawn from a single city, and the reviews were published in a single language. Although Moscow is a cosmopolitan, international destination, our data could reflect a bias of travelers who visit that city.

Third, a consequence of Moscow's international status is that our data contained a sizeable number of reviews that were written in languages other than English. Since we analyzed only the English-language reviews, it may be that the total corpus would demonstrate different results. Text analysis across multiple languages presents methodological difficulties. However, when those issues are overcome, online reviews will potentially yield insights about cultural effects that can further aid hotel managers in improving their customer experiences.

Fourth, our study data were compiled over a relatively brief time span. We see the likely benefit of an analysis that takes into account how reviews have evolved over time. Such a longitudinal study could investigate the role of influential reviewers, or quantify how online customer feedback affects the financial health of hotels when the online reviews are combined with financial performance data. Another limitation to be overcome involves developing a customized dictionary to better categorize the sentiment in hotel reviews. Since the dictionaries we used were generated for purposes other than hospitality operations, they may have an inherent bias. Building a sentiment dictionary that is specific to the hospitality industry would allow more precise analysis.

Finally, even though topic modeling is a powerful tool for text analytics, it's also a bit *ad hoc*, a situation that represents a final potential study limitation. The topics we arrived at (and the words included in each topic) were automatically generated and seemed to make sense at face value, but could be further adjusted and refined in consultation with subject matter experts. Developing text analysis methods that incorporate domain knowledge could be a fruitful path to generating deeper managerial insights. In sum, it's clear that the industry should go beyond numerical ratings and pay attention to the text found in reviews. Guests' true feelings are found in those comments—particularly if they write a lengthy review that focuses tightly on just a few issues. ■

APPENDIX

A Partial List of Companies That Provide Text Analytics

Company	Sentiment	Topic Modeling
ReviewPro resources.reviewpro.com/guides/review-analytics-hotel-investors	Yes—Automated	Yes—Automated with predefined dictionary
CzarMetrics www.czarmetrics.com	Yes—Manual	Yes—Manual
Fishbowl Analytics www.fishbowl.com/analytics/	Yes—Automated	No
Where2GetIt www.where2getit.com/solutions/local-seach-suite/review-management	Yes—Automated	No
Brandify www.brandify.com/HowBrandifyWorks/	Yes—Automated	No
ReviewTrackers www.reviewtrackers.com/tour	No	No
Hoxell www.hoxell.com	No	No

Center for Hospitality Research
Publication Index
chr.cornell.edu

2016 Reports

2015 Reports

2015 Tools

Cornell Hospitality Report
Vol. 16, No. 2 (February 2016)

Cornell Hospitality Report is produced for the benefit of the hospitality industry by The Center for Hospitality Research at Cornell University.

Michael C. Sturman, Director
Carol Zhe, Program Manager
Glenn Withiam, Executive Editor
Alfonso Gonzalez, Executive Director of Marketing and Communications

Center for Hospitality Research
Cornell University
School of Hotel Administration
537 Statler Hall
Ithaca, NY 14853

607-255-9780
chr. cornell.edu