12-15-2019

Trust-based content filtering: Investigating the association between assurance seals, source expertise, and topics of online product reviews

Bartlomiej Hanus

Benjamin George

Trust-Based Content Filtering: Investigating the Association between Assurance Seals, Source Expertise, and Topics of Online Product Reviews

Bartlomiej Hanus[1]
Beacom School of Business, University of South Dakota,
Vermillion, SD, USA

Benjamin George
Beacom School of Business, University of South Dakota,
Vermillion, SD, USA

ABSTRACT

Online product reviews are a significant component affecting transactions in business-to-consumer (B2C) e-commerce. The sheer volume of online reviews makes it virtually impossible for buyers to systematically process all reviews available. Drawing on the elaboration likelihood model (ELM) and web assurance seals (AS) literature, we investigate the association between two trust-building proxies included in reviews: verified buyer flag (VBF) and reviewer's technical understanding (TU), and topics discussed in online product reviews. Our results indicate that both VBF and TU affect review content. From a practical perspective, we provide a means of content filtering that can be implemented at a recommender system level to reduce information overload prospective buyers are subjected to. From a theoretical perspective our results indicate there is an identifiable shift that has occurred in the e-commerce environment. More specifically, the evolution of the web has brought elements of consumer-to-consumer (C2C) interactions into the space typically reserved for B2C landscape, where sellers also act as intermediaries facilitating information exchange between buyers.

Keywords: trust, online reviews, topic modeling, source credibility, assurance seals, security software, antivirus.

---

[1] Corresponding author. Bart.Hanus@usd.edu + 1 605 658 6542

INTRODUCTION

Online reviews are the most influential factor affecting purchasing decisions of online shoppers and is more influential than advertising or friends and family recommendations (BigCommerce 2017). With reviews being different in terms of content, emotional load, or intended audience, and their increasing volume, it becomes increasingly difficult to distinguish the signal from the noise, as most buyers will not read all available reviews (Shrestha 2018). So, the question remains which reviews are to be trusted.

Sellers often implement measures intending to increase trustworthiness of reviews posted on their portals. In this study, we investigate two of such measures: verified buyer flag (VBF) and self-reported reviewer level of technical understanding (TU). VBF symbolizes a seller issued AS certifying if an author of a review purchased the product, that way increasing legitimacy of a review. TU epitomizes reviewer's self-reported level of technical expertise with intention to legitimize reviews based on the assumption that subject matter experts can be trusted more than novices. The primary goal of the present study is to examine if there exists any association between the two trust-building proxies and topics discussed in product reviews.

We collected online reviews of security software from Newegg.com, a popular online seller. We then extracted the topics of the reviews using Latent Dirichlet Allocation (LDA) and performed loglinear analysis to assess the association between the three variables: topics, VBF, and TU. Our results indicate that both trust-related variables affect the topics being discussed in the reviews.

THEORETICAL BACKGROUND

AS have roots in research founded on the premise that individuals tend to place their trust in settings that increase their feeling of safety (Mcknight et al. 2004). Their purpose is to increase

buyers' trust in the seller's platform given the uncertainty of the e-commerce environment. AS are usually associated with using third-party services, which act as certifying entities. The organizations typically have a set of standards developed, with which online sellers must comply (and periodically verify) to obtain a specific institution's AS to be used on the seller's website. The purpose of these seals is to instill trust in buyers reassuring them that the seller's privacy practices, safety of transactions and records have met the highest standards (Kim et al. 2008).

Several studies confirm AS effectiveness with their positive effect on willingness to buy (Kovar et al. 2000; Mauldin and Arunachalam 2002; Nöteberg et al. 2003), while simultaneously generating higher expectations towards transactions (Rifon et al. 2005). Similar findings were reported by Park et al. (2010) who found that customers elicited higher expectations towards vendors with seals, but in case of their absence they would shift their focus towards sellers' service performance. Additionally, Nöteberg et al. (2003) noticed that presence of AS decreases buyers' privacy concerns, regardless of who the provider of the seal was. In contrast, Odom et al. (2002) discovered seals issued by different providers addressed different types of customer concerns other than privacy (e.g., security, legitimacy, documentation, quality, etc.). AS that perform multiple functions are not more efficient than those focusing on one only (Hu et al. 2010). There are also studies that find no effects of AS on trust (Kimery and McCord 2002; Mcknight et al. 2004; Wang et al. 2004). The exact reasons for these inconclusive findings are not fully understood.

We focus on a seller issued VBF that is placed next to each individual review, so that prospective buyers can easily notice it. We propose that the role of VBF is the same as any other AS, including those that are issued by third parties – to instill trust. However, this time, the issuer of the seal is shifted, whereas the recipient remains unchanged. Concurrently, seller-issued AS

(like VBF) will only work as intended, if the seller has a reputation of being solid, which – paradoxically – may often be influenced by AS issued by third parties to the seller.

Initial trust in online reviews relates to cases where buyers do not have reliable means to determine if reviews they read can be depended upon (McKnight et al. 2002). It is formed in short period of time during the first encounter with the trustee (McKnight et al. 1998) and may affect purchasing decision. The level of trust that a trustee can raise in the trustor is tied to the effectiveness of their persuasion. Drawing on the elaboration likelihood model (Chen and Chaiken 1999; Petty and Cacioppo 1986), individuals can process information either in a systematic or heuristic mode, both of which can affect their attitudes. The former requires careful evaluation of the arguments presented, while the latter often depends on peripheral cues which can affect the attitude in the absence of systematic information processing.

The major determinants of persuasion include personal relevance, argument quality, and source expertise (Pornpitakpan 2004; Yi et al. 2013). Petty & Cacioppo (1986) determined that persuasion methods are more effective when arguments presented were of high quality and their source demonstrated high levels of expertise. We propose that when buyers are exposed to a large volume of reviews, they will first engage their peripheral cue-based heuristic processing to identify reviews relevant to them, and then read those to obtain further information. One type of peripheral cues based on source credibility is reviewer's TU. Based on prior studies, prospective buyers will trust reviews written by experts more than they will those written by novices. From this perspective, both VBF and the reviewer's TU can be described as a form of peripheral cues attached to an online review, whose underlying goal is to increase trustworthiness of reviews. The goal of the present study is to investigate the interplay between these two cues and the argument presented in a body of a review.

METHODOLOGY

We collected online reviews from Newegg.com, a seller with long market presence (founded in 2001) and high ratings among buyers (ResellerRatings 2019) and third-parties (Better Business Bureau 2019). We focused on antivirus and internet security software. We extracted product reviews using a custom-built scraper, yielding a sample of 6622 unique reviews. For each, we collected product id, product name, manufacturer, reviewer name/id, TU, length of ownership, and VBF) and review components (i.e., the body, including pros, cons, and other thoughts sections, product rating assigned, and review date). VBF is represented with a Boolean variable. TU is reported optionally by reviewers and described on a five-point scale, ranging between low and high or set to unknown (if not reported). Pros, cons, and other thoughts were merged into single documents. We then conducted additional text preprocessing: tokenization, lower-case transformation, and stop words removal. Next, we filtered out tokens shorter than three and longer than thirty characters and then generated bigrams. Lastly, we pruned tokens appearing in less than one percent or in more than ten percent of documents. We then applied LDA, a method which assumes a fixed number of topics with each document in a corpus representing a combination of the above topics. LDA operates by creating a posterior structure representing a summary that can be used to explore the underlying corpus of text. It generates relative topic proportions and topic assignment for each document (Blei et al. 2010).

Determining the "optimal" number of topics with LDA is not well-documented. While several alternatives exist (Arun et al. 2010; Cao et al. 2009; Griffiths and Steyvers 2004), we took an approach based on LDA's close links to principal component analysis (PCA) (Blei et al. 2010). We first ran LDA with a large number of topics and using the scree test we determined that the number of topics extracted was six. To check for association between VBF, TU, and

topics we implemented loglinear analysis. LDA in addition to generating topic proportions also provided us with topic assignment for each document. Thus, each review had one primary topic assigned to it (used as a categorical variable in our model). Finally, we calculated frequencies for each combination of topic assignment, TU, and VBF.

DATA ANALYSIS

Table 1 includes typical summary diagnostics for the six topics extracted using LDA (McCallum 2002). The tokens row provides the number of word tokens allotted to each topic, giving the proportion of the corpus assigned to each topic. These range between 29,202 for topic 2 and 242,244 for topic 1. These numbers should not be small or large relatively to other topics. Document entropy ranges between 6.8578 (topic 5) and 8.3442 (topic 1) and does not indicate any reasons for concern. Some variation is present in document coherence (if terms in a topic tend to co-occur together). Large negative values represent less coherent topic, while those closer to zero indicate more coherent one. Topics 1 and 0 are the most coherent, while topics 2 and 5 are the least. In terms of uniform and corpus distributions, topic 1 stands out the most. Corpus distribution is correlated with the number of tokens with the assumption that topic that incorporates a relatively large fraction of tokens in the document collection will be more like the overall corpus distribution (i.e., topic 1). Based on the effective number of words metric, topics 2, 3, and 5 are characterized by high specificity, while topics 1, 0, and 4 are its lower range. Token/Document discrepancy examines how many times a given term occurs in a specific topic in comparison to the number of documents that the term occurs in as part of the topic. This metric is somewhat related to coherence, with topic 5 demonstrating the highest discrepancy. Rank 1 documents metric evaluates the number the instances where a specific topic is the most frequent topic in the document. In other words, it indicates whether topics are "main" topics

(higher values) or "background" topics (lower values). As such, topics 1, 0, and 4 appear to be the main topics, while topics 2, 5, and 3 can be of lesser significance (they represent a distinct discourse in a document but are not necessarily the focus of that document).

Table 1. LDA Topic Model Diagnostics

Topic	Topic_0	Topic_1	Topic_2	Topic_3	Topic_4	Topic_5
Coherence	-66.68	-64.20	-154.76	-118.46	-88.81	-141.60
Rank 1 Documents	0.297	0.745	0.034	0.05	0.211	0.039
Average Word Length	6.0	5.8	8.8	3.7	6.3	7.5
Exclusivity	0.8576	0.8346	0.7264	0.6962	0.7751	0.7424
Document Entropy	6.8962	8.3442	6.9003	6.8438	7.2587	6.8578
Effective Number of Words	936.5480	727.5651	9913.2408	7988.2533	1146.1534	5546.9356
Tokens	55465	242244	29202	30152	59765	31194
Uniform Dist.	2.8254	3.4847	2.1151	2.1573	2.6645	2.1969
Corpus Dist.	1.5839	0.4691	2.4523	2.4091	1.5613	2.3664
Token/Document Discrepancy	0.0023	0.0013	0.0060	0.0002	0.0019	0.0105
Allocation Count	0.3511	0.8750	0.1743	0.1859	0.3354	0.1870

Table 2 includes the list of top ten terms for each of the six topics identified. We have also provided a label for each topic that describes the area that the topic is discussing. We labeled topic 0 "Data recovery". The reviews falling in this category primarily discuss data backup solutions, which are usually bundled with most modern consumer-oriented computer security suites. One of the terms – "Acronis" refers a software vendor whose primary area of expertise is data protection. Reviews assigned to topic 1 "General functionality" usually describe the users' general experience with the software package and the main purpose for which it was designed – keeping the system free of viruses and malware. The product name "Norton" is also one of the major terms characterizing topic 1. Nearly 22% of all reviews were about Symantec products. Topic 1 is also the main topic in our collection based on rank 1 metric, which confirms our expectations as the reviews falling in this category describe the primary functionality of this type of software. Topic 2 "Virus detection" involves the reviewers' experience with false alarm detection, a problem that has been nagging the industry for years. This primarily occurs when

detection signatures or detection algorithm employed by an antivirus package when scanning the system are not finely tuned. Topic 3 "Hardware resource" discusses what hardware software was used on, running scan at boot time, and the application's impact on system performance. It is usually one of the metrics that antivirus software is evaluated on, with the expectation that it should be running in the background without consuming a significant proportion of the system's resources. Topic 4 "Product support" includes items that are related to either the vendor support or the seller support from Newegg. Finally, topic 5 "Configurability" includes the discussion of additional functionality of a given package. For instance, these could include setting up parental controls, optimizing and cleaning Windows registry, setting up firewall, etc. Some of these items might be configured right out of the box, some may require user intervention first.

Table 2. Top 10 Terms per Topic

Topic	Label	Top 10 Terms
0	Disk recovery	drive, backup, acronis, image, disk, restore, windows, software, ghost, drives
1	General functionality	software, norton, product, computer, use, system, virus, good, version, free
2	Virus detection	false, positives, false_positives, nav, infections, response, infected, manufacturer, comparatives, please
3	Hardware resources	core, cpu, ram, ghz, dual, amd, gig, boot, blue, laptop
4	Product support	support, product, rebate, customer, tech, newegg, tech_support, get, service, key
5	Configurability	webroot, devices, settings, registry, parental, options, device, unlimited, control, firewall

When calculating frequencies of each combination of VBF, TU, and Topic, we filtered the cells for which the count was less than five and did not include them in subsequent analyses to avoid potential threats to validity (Allison 2012). We then fitted a loglinear model to evaluate the association between the three categorical variables. We initially fitted a saturated model only to find out that the 3-way interaction between the categorical variables was not significant. We then fitted a more parsimonious, homogenous association model with the highest order interaction removed. Table 3 includes the reduced model's goodness of fit statistics. Both deviance and Pearson chi-square have values to df ratios close to 1, indicating good fit (Allison

2012). Homogeneous association model assumes that all partial odds are resulting from two-way interaction terms. The analysis of those is presented in Table 4. All three two-way interactions are significant. To further investigate the current model's fit, we also analyzed its residuals. The largest value was less than 2.4, providing sufficient evidence for our model's good fit.

Table 3. Goodness of Fit Assessment – Homogeneous Association Loglinear Model

Criterion	DF	Value	Value/DF
Deviance	11	11.7117	1.0647
Scaled Deviance	11	11.7117	1.0647
Pearson Chi-Square	11	11.9802	1.0891
Scaled Pearson X2	11	11.9802	1.0891
Log Likelihood		27917.5866	
Full Log Likelihood		-133.0179	
AIC (smaller is better)		340.0358	
AICC (smaller is better)		621.2358	
BIC (smaller is better)		409.2702	

Table 4. Likelihood Ratio Statistics for Type 3 Analysis

Source	DF	Chi-Square	Pr > ChiSq
VBF	1	12.22	0.0005
TU	5	543.60	<.0001
Topic	5	1223.53	<.0001
VBF * TU	5	110.45	<.0001
VBF * Topic	5	11.45	0.0431
TU * Topic	15	125.21	<.0001

Detailed list of maximum likelihood parameter estimates is included in Appendix A. If there are significant interaction terms found in the model, then lower-order terms are omitted from interpretation, as well as the intercept which is a normalizing constant. The main effects for VBF, TU, and topic provide us with little meaningful insight and are difficult to interpret. The estimates included in Appendix A represent conditional log odds ratios. Those that are of interest to us are the estimates representing two-way associations identified as significant in Appendix A (i.e., VBF * TU and TU * Topic).

First, reviewers who are verified buyers are less likely to report their TU levels with extreme cases of high, somewhat high, and average TU levels. These negative associations are independent of the topics discussed in the body of the reviews. Second, independent of VBF, there are associations between their levels of TU and topics in their reviews, including the following: (1) negative association between low TU and topic 1, (2) positive association between "average" TU and topics 0 and 1, (3) positive association between "somewhat high" TU and topics 0, 1, and 4, (4) positive association between "high" TU and topics 0 and 1. Finally, independent of TU, as far as the VBF and topic association is concerned, verified buyers tend to discuss topic 1 and, perhaps topics 4 and 0 more frequently than other topics.

DISCUSSION

The six-topic solution represents most of the typical applications of antivirus and security software designed for consumer market. Consumer software should work for the most part right out-of-the-box, while still offering some degree of customization. We believe that our six-topic solution reflects these properties closely. These range from general functionality (topic 1) that describes the overall product experience and usage without going into more specific details, through the service quality component (topic 4), to more detailed aspects of software. The service quality component includes items related to the seller (i.e., Newegg) and those associated with software vendors. The remaining four topics describe more specific areas. For instance, topic 2 "virus detection" describes the accuracy of its malware detection engine. These should be as accurate as possible. Most vendors promise high detection rates, but 100% is almost never guaranteed. In fact, software warranties offer little customer protection and limit vendor liability (Riedy and Hanus 2017). For consumer-oriented security suites, 99.5% detection rate is considered a minimum threshold (Virus Bulletin 2019), but is more of an expectation than a legal

obligation. Regarding detection errors, most usually are false positives, while false negatives are fewer in numbers. From end-user perspective it can an inconvenience if healthy files are quarantined. Topic 3 "hardware resources" captures portions of reviews that are related to overall performance of a given product. While historically some software offered only on-demand scanning, modern solutions offer real-time protection service running in the background. It can be invoked for events like file copying, installing or removing applications, archiving files, browsing the web and downloading files, etc. While these may not be resource intensive, adding an additional layer of protection can impact system's performance. Thus, security software should leave a small footprint on computer's memory and CPU utilization. Reviewers also discuss configurability of software (topic 5). It may be related to the number of concurrent installations and whether they can be centrally managed. Most vendors have multiple licensing options available to cover multiple devices that consumers have. More sophisticated options may include setting up firewall or parental controls, or cleaning Windows registry form orphaned entries, allowing users to finetune the software. Lastly, topic 0 "disk recovery" is related to disk backup and recovery that can help avoiding data loss in case of malware infection or ransomware. Historically data backup packages were sold separately from antivirus software, whereas these days they tend to be bundled together. Overall, the six-topic solution offers a comprehensive view of areas buyers should consider when shopping of antivirus software.

As shown in Appendix A, the parameter estimates for interaction terms describing the association between VBF and TU levels shed additional light into the interdependencies between these two trust proxies. Verified buyers are less likely to report high TU in comparison to non-verified buyers, and to a lesser degree for somewhat high and average TU levels. It appears that verified buyers may simply not bother about indicating their TU as opposed to their non-verified

counterparts. A plausible explanation would be that VBF carries enough significance in the eyes of reviewers, so they assume no additional source credibility measures are to instill trust in prospective buyers. However, it is not guaranteed that every verified buyer will know how to use the product as intended and that they will share accurate information that could help prospective buyers make a purchase decision. This potentially brings the role of review helpfulness into the picture and could be further evaluated by future studies. The data we collected were insufficient with regards to this metric. It appears that non-verified buyers may require a substitute in the absence of VBF in order to substantiate the claims in their reviews. When reported, most reviewers tend to evaluate their TU as either high or somewhat high, potentially raising concerns on the usefulness of the metric, given that verified buyers tend to give it less consideration.

The interaction between VBF and topics sheds additional light on the association between the VBF and TU. Verified buyers tend to mostly focus on the software's general functionality (topic 1), product support (topic 4) and data recovery (topic 0), and tend to be less concerned with more "technical" issues like virus detection (topic 2), hardware resources (topic 3), or configurability (topic 5). Just as they were not concerned about indicating their TU, verified buyers are inclined to provide more general information in their reviews, without going to additional details. One possible explanation is that they want to communicate to prospective buyers if the software works and what assistance, if any, can they expect from vendors. In contrast, it appears that non-verified buyers are more interested in more technical details in their reviews, which – to an extent – is in line with their tendency to report high TU levels. These two items together create an impression that non-verified buyers might simply be bragging about their expertise and that way try to convince prospective buyers their reviews should be trusted.

Lastly, interesting insights can be gained from the association between TU and topics. Individuals who report low TU are less likely to discuss topic 1 than topic 4. This could be simply because, in comparison to more experienced people, these users lack skills to operate the software and, thus, heavily rely on vendor support to get software installed and set up. This observation, however, does not apply to individuals with "somewhat low" TU, potentially indicating that these two groups are separated by the latter's ability to get their software up and running without any external assistance. We also found that individuals reporting the two lowest levels of TU do not discuss the other four topics. Those who report average, somewhat high, or high TU levels are much more diverse in that regards. At the same time, the members of these three groups are more likely to discuss issues related to data recovery than they are to mention any other topics.

CONCLUSION

Overall, our results indicate there is association between VBF, TU, and topics of antivirus software reviews. Verified buyers tend to be more straight-to-the-point in their reviews and often avoid going into highly technical details behind each software package. Users with lower levels of technical understanding focus on product support and general software functionality. Their more experienced counterparts are broader in their deliberations, often touching more detailed subjects (i.e., data recovery). Our results also indicate that VBF and TU tend to be used, but not necessarily in unison. Reviewers who are verified buyers are more likely not to report TU, while non-verified buyers will attempt to substitute the absence of VBF with higher levels of TU.

Our study extends the current body of knowledge by investigating trust-building proxies in a previously not discussed setting. Typically, as extant literature demonstrates, AS are issued by third parties to build-up trust for sellers. We, on the other hand, investigated the use of AS

issued by sellers to reviewers (i.e., VBF). In addition, we find that the use of TU as a source of credibility occurs in the form of customer interaction historically kept for C2C, now deployed in a B2C with sellers acting as intermediaries. In e-commerce, these instruments are important mechanisms allowing differentiating between online reviews. Future studies could investigate which reviews (based on topics) are determinants of purchasing intentions.

From practical perspective, our results provide sellers with an outlet for improved content filtering systems available on their platform to prospective buyers that would be customized to reflect their level of TU. Sellers could allow searching and filtering for those reviews containing topics relevant to prospective buyers, allowing them to find the reviews quicker and, more importantly, save their time by not having to read the other reviews.

REFERENCES

Allison, P. D. 2012. *Logistic Regression Using SAS: Theory and Application, Second Edition*, (2nd ed.), Cary, NC, USA: SAS Institute.

Arun, R., Suresh, V., Veni Madhavan, C. E., and Narasimha Murthy, M. N. 2010. "On Finding the Natural Number of Topics with Latent Dirichlet Allocation: Some Observations," in *Advances in Knowledge Discovery and Data Mining*, M. J. Zaki, J. X. Yu, B. Ravindran, and V. Pudi (eds.), Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 391–402.

Better Business Bureau. 2019. "Newegg, Inc. | Better Business Bureau® Profile." (https://www.bbb.org/us/ca/whittier/profile/computer-parts/newegg-inc-1216-13146135, accessed May 27, 2019).

BigCommerce. 2017. "Omni-Channel Retail in 2017 What Brands Need to Know and Modern Consumer Shopping Habits." (https://grow.bigcommerce.com/rs/695-JJT-333/images/the-omni-channel-selling-guide.pdf).

Blei, D. M., Carin, L., and Dunson, D. 2010. "Probabilistic Topic Models A Focus on Graphical Model Design and Applications to Document and Image Analysis," *IEEE Signal Processing Magazine* (27:6), pp. 55–65. (https://doi.org/10.1038/mp.2011.182.doi).

Cao, J., Xia, T., Li, J., Zhang, Y., and Tang, S. 2009. "A Density-Based Method for Adaptive LDA Model Selection," *Neurocomputing* (72:7–9), Elsevier, pp. 1775–1781. (https://doi.org/10.1016/J.NEUCOM.2008.06.011).

Chen, S., and Chaiken, S. 1999. "The Heuristic-Systematic Model in Its Broader Context," *Dual-Process Theories in Social Psychology*, pp. 73–96.

Griffiths, T. L., and Steyvers, M. 2004. "Finding Scientific Topics," *Proceedings of the National Academy of Sciences* (101:suppl 1), 5228 LP – 5235. (https://doi.org/10.1073/pnas.0307752101).

Hu, X., Wu, G., Wu, Y., and Zhang, H. 2010. "The Effects of Web Assurance Seals on Consumers' Initial Trust in an Online Vendor: A Functional Perspective," *Decision Support Systems* (48:2), pp. 407–418. (https://doi.org/10.1016/j.dss.2009.10.004).

Kim, D. J., Steinfield, C., and Lai, Y.-J. 2008. "Revisiting the Role of Web Assurance Seals in Business-to-Consumer Electronic Commerce," *Decision Support Systems* (44:4), pp. 1000–1015.

Kimery, K., and McCord, M. 2002. "Third-Party Assurances: The Road to Trust in Online Retailing," in *Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS'02)-Volume 7 - Volume 7*, HICSS '02, Washington, DC, USA: IEEE Computer Society, pp. 175--.

Kovar, S. E., Burke, K. G., and Kovar, B. R. 2000. "Consumer Responses to the CPA WEBTRUST™ Assurance," *Journal of Information Systems* (14:1), pp. 17–35. (https://doi.org/10.2308/jis.2000.14.1.17).

Mauldin, E., and Arunachalam, V. 2002. "An Experimental Examination of Alternative Forms of Web Assurance for Business☐to☐Consumer E☐Commerce," *Journal of Information Systems* (16:s-1), pp. 33–54. (https://doi.org/10.2308/jis.2002.16.s-1.33).

McCallum, A. K. 2002. "MALLET: A Machine Learning for Language Toolkit." (http://mallet.cs.umass.edu).

McKnight, D. H., Choudhury, V., and Kacmar, C. 2002. "Developing and Validating Trust Measures for E-Commerce: An Integrative Typology," *INFORMATION SYSTEMS RESEARCH* (13:3), INFORMS: Institute for Operations Research, pp. 334–359.

McKnight, D. H., Cummings, L. L., and Chervany, N. L. 1998. "Initial Trust Formation in New Organizational Relationships," *Academy of Management Review* (23:3), pp. 473–490. (https://doi.org/10.5465/amr.1998.926622).

Mcknight, H. D., Kacmar, C. J., and Choudhury, V. 2004. "Shifting Factors and the Ineffectiveness of Third Party Assurance Seals : A Two-Stage S," *Electronic Markets* (14:3), pp. 1–15.

Nöteberg, A., Christiaanse, E., and Wallage, P. 2003. "Consumer Trust in Electronic Channels," *EService Journal* (2:2), pp. 46–67. (https://doi.org/10.1353/esj.2003.0014).

Odom, M. D., Kumar, A., and Saunders, L. 2002. "Web Assurance Seals: How and Why They Influence Consumers' Decisions," *Journal of Information Systems* (16:2), pp. 231–250.

Park, I., Bhatnagar, A., and Rao, H. R. 2010. "Assurance Seals, on-Line Customer Satisfaction, and Repurchase Intention," *International Journal of Electronic Commerce* (14:3), pp. 11–34. (https://doi.org/10.2753/JEC1086-4415140302).

Petty, R. E., and Cacioppo, J. T. 1986. "The Elaboration Likelihood Model of Persuasion," in *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*, New York, NY: Springer New York, pp. 1–24. (https://doi.org/10.1007/978-1-4612-4964-1_1).

Pornpitakpan, C. 2004. "The Persuasiveness of Source Credibility: A Critical Review of Five Decades' Evidence," *Journal of Applied Social Psychology* (34:2), pp. 243–281.

ResellerRatings. 2019. "Newegg.Com (USA) Reviews | 48,870 Reviews of Newegg.Com | ResellerRatings." (https://www.resellerratings.com/store/Newegg, accessed May 27, 2019).

Riedy, M. K., and Hanus, B. 2017. "It Is Just Unfair Using Trade Laws to out Security Software Vulnerabilities," *Loy. U. Chi. LJ* (48:4), pp. 1099–1134.

Rifon, N. J., Larose, R., and Choi, S. M. 2005. "Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures," *Journal of Consumer Affairs* (39:2), pp. 339–362. (https://doi.org/10.1111/j.1745-6606.2005.00018.x).

Shrestha, K. 2018. "50 Important Stats You Need to Know About Online Reviews [Infographic]," *Vendasta Blog*. (https://www.vendasta.com/blog/50-stats-you-need-to-know-about-online-reviews, accessed April 28, 2018).

Virus Bulletin. 2019. "Virus Bulletin's Renowned VB100 Certification Programme." (https://www.virusbulletin.com/testing/vb100, accessed July 19, 2019).

Wang, S., Beatty, S. E., and Foxx, W. 2004. "Signaling the Trustworthiness of Small Online Retailers," *Journal of Interactive Marketing* (18:1), pp. 53–69. (https://doi.org/https://doi.org/10.1002/dir.10071).

Yi, M. Y., Yoon, J. J., Davis, J. M., and Lee, T. 2013. "Untangling the Antecedents of Initial Trust in Web-Based Health Information: The Roles of Argument Quality, Source Expertise, and User Perceptions of Information Quality and Risk," *Decision Support Systems* (55:1), pp. 284–295. (https://doi.org/10.1016/j.dss.2013.01.029).

APPENDIX A – ANALYSIS OF MAXIMUM LIKELIHOOD PARAMETER

ESTIMATES

Parameter			Estimate	Standard Error	Likelihood Ratio 95% Confidence Limits		Wald Chi-Square	Pr > Chi Sq
Intercept			1.3712	0.3330	0.6720	1.9831	16.95	<.0001
VBF			0.8327	0.2645	0.3146	1.3549	9.91	0.0016
TU	Low		-1.0376	0.5516	-2.1677	0.0111	3.54	0.0600
TU	Somewhat Low		-1.9776	0.5952	-3.2444	-0.8821	11.04	0.0009
TU	Average		0.1611	0.4313	-0.7089	1.0038	0.14	0.7088
TU	Somewhat High		0.5199	0.4035	-0.2790	1.3203	1.66	0.1976
TU	High		1.5630	0.3429	0.9146	2.2691	20.78	<.0001
TU	Unknown		0.0000	0.0000	0.0000	0.0000	.	.
Topic	Topic_0		-1.0792	0.5617	-2.2721	-0.0263	3.69	0.0547
Topic	Topic_1		2.9500	0.3317	2.3371	3.6441	79.11	<.0001
Topic	Topic_2		0.2533	0.4640	-0.6621	1.1689	0.30	0.5852
Topic	Topic_3		0.0669	0.2807	-0.4827	0.6217	0.06	0.8115
Topic	Topic_4		1.4448	0.3522	0.7862	2.1739	16.83	<.0001
Topic	Topic_5		0.0000	0.0000	0.0000	0.0000	.	.
VBF*TU	Low		-0.7123	0.4526	-1.5766	0.2229	2.48	0.1156
VBF*TU	Somewhat Low		-0.5607	0.3818	-1.2875	0.2236	2.16	0.1420
VBF*TU	Average		-0.6839	0.1355	-0.9527	-0.4210	25.47	<.0001
VBF*TU	Somewhat High		-0.8716	0.1240	-1.1186	-0.6322	49.43	<.0001
VBF*TU	High		-1.1140	0.1208	-1.3551	-0.8811	85.07	<.0001
VBF*TU	Unknown		0.0000	0.0000	0.0000	0.0000	.	.
VBF*Topic	Topic_0		0.4848	0.2599	-0.0263	0.9962	3.48	0.0622
VBF*Topic	Topic_1		0.5081	0.2481	0.0198	0.9964	4.19	0.0405
VBF*Topic	Topic_2		-0.2684	0.3907	-1.0406	0.4955	0.47	0.4921
VBF*Topic	Topic_3		0.2706	0.3319	-0.3802	0.9234	0.66	0.4149
VBF*Topic	Topic_4		0.4990	0.2618	-0.0158	1.0140	3.63	0.0567
VBF*Topic	Topic_5		0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Low	Topic_1	-1.2041	0.4162	-2.0147	-0.3657	8.37	0.0038
TU*Topic	Low	Topic_4	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Somewhat Low	Topic_1	0.0543	0.5073	-0.8635	1.1667	0.01	0.9147
TU*Topic	Somewhat Low	Topic_4	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Average	Topic_0	2.9185	0.6267	1.7474	4.2360	21.69	<.0001
TU*Topic	Average	Topic_1	0.8841	0.4279	0.0489	1.7486	4.27	0.0388
TU*Topic	Average	Topic_3	0.1842	0.4725	-0.7339	1.1333	0.15	0.6967
TU*Topic	Average	Topic_4	0.4316	0.4502	-0.4472	1.3372	0.92	0.3377
TU*Topic	Average	Topic_5	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Somewhat High	Topic_0	3.5255	0.6035	2.3976	4.7980	34.13	<.0001
TU*Topic	Somewhat High	Topic_1	1.4481	0.4002	0.6550	2.2417	13.09	0.0003
TU*Topic	Somewhat High	Topic_2	-0.0453	0.6383	-1.3445	1.1829	0.01	0.9434

Parameter			Estimate	Standard Error	Likelihood Ratio 95% Confidence Limits		Wald Chi-Square	Pr > Chi Sq
TU*Topic	Somewhat High	Topic_3	0.5222	0.4119	-0.2732	1.3502	1.61	0.2049
TU*Topic	Somewhat High	Topic_4	0.8267	0.4188	-0.0015	1.6560	3.90	0.0484
TU*Topic	Somewhat High	Topic_5	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	High	Topic_0	3.1380	0.5636	2.0855	4.3367	31.00	<.0001
TU*Topic	High	Topic_1	0.8409	0.3392	0.1426	1.4835	6.14	0.0132
TU*Topic	High	Topic_2	-0.1465	0.4692	-1.0742	0.7758	0.10	0.7548
TU*Topic	High	Topic_3	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	High	Topic_4	0.1410	0.3599	-0.5935	0.8272	0.15	0.6952
TU*Topic	High	Topic_5	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Unknown	Topic_0	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Unknown	Topic_1	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Unknown	Topic_2	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Unknown	Topic_4	0.0000	0.0000	0.0000	0.0000	.	.
TU*Topic	Unknown	Topic_5	0.0000	0.0000	0.0000	0.0000	.	.
Scale			1.0000	0.0000	1.0000	1.0000		