Title
Detecting fake-review buyers using network structure: Direct evidence from Amazon

Journal
Proceedings of the National Academy of Sciences of the United States of America, 119(47)

ISSN
0027-8424

Authors
He, Sherry
Hollenbeck, Brett
Overgoor, Gijs
et al.

Publication Date
2022-11-22

DOI
10.1073/pnas.2211932119

Peer reviewed

PNAS RESEARCH ARTICLE | SOCIAL SCIENCES

Detecting fake-review buyers using network structure: Direct evidence from Amazon

Sherry He[a], Brett Hollenbeck[a,1], Gijs Overgoor[b], Davide Proserpio[c], and Ali Tosyali[b]

Edited by Avi Goldfarb, University of Toronto, Canada; received July 13, 2022; accepted October 10, 2022 by Editorial Board Member Mark Granovetter

Online reviews significantly impact consumers' decision-making process and firms' economic outcomes and are widely seen as crucial to the success of online markets. Firms, therefore, have a strong incentive to manipulate ratings using fake reviews. This presents a problem that academic researchers have tried to solve for over two decades and on which platforms expend a large amount of resources. Nevertheless, the prevalence of fake reviews is arguably higher than ever. To combat this, we collect a dataset of reviews for thousands of Amazon products and develop a general and highly accurate method for detecting fake reviews. A unique difference between previous datasets and ours is that we directly observe which sellers buy fake reviews. Thus, while prior research has trained models using laboratory-generated reviews or proxies for fake reviews, we are able to train a model using actual fake reviews. We show that products that buy fake reviews are highly clustered in the product reviewer network. Therefore, features constructed from this network are highly predictive of which products buy fake reviews. We show that our network-based approach is also successful at detecting fake review buyers even without ground truth data, as unsupervised clustering methods can accurately identify fake review buyers by identifying clusters of products that are closely connected in the network. While text or metadata can be manipulated to evade detection, network-based features are more costly to manipulate because these features result directly from the inherent limitations of buying reviews from online review marketplaces, making our detection approach more robust to manipulation.

online reviews | networks | machine learning | text analysis

Online reviews have a significant impact on consumer purchase decisions and are widely seen as crucial to the success of online markets (1, 2). Review and rating systems allow buyers and sellers to develop credible reputations in settings that are otherwise mostly anonymous. Because these reputations are crucial for seller outcomes, sellers have a large incentive to manipulate their ratings and inflate their reputation. As a result, online review platforms like Amazon, Yelp, or Tripadvisor have struggled since their inception with the problem of sellers manipulating their ratings with fake reviews. Rating manipulation can potentially cause buyers to buy from lower-quality sellers than they otherwise would, allow sellers to charge higher prices than if their true reputation was observed, and lower trust in reviews and review platforms, making it difficult for high-quality and honest sellers to compete. There is growing empirical evidence that fake reviews harm consumers (3). In addition to violating the platform's policies, these practices are the subject of ongoing investigations by the Federal Trade Commission, the United Kingdom Competition and Markets Authority, and other regulators.

Despite the vast amount of academic research over the past two decades (see ref. 4 for an extensive review) and the large amounts of time, effort, and money invested by online platforms to detect and remove fake reviews, they are nevertheless as prevalent as ever. Recent studies have found that millions of products on Amazon are using fake reviews, a large share of all reviews are fake (3), and consumers express very low levels of trust in online reviews as a result (5).

Most research on the challenge of fake-review detection relies on machine learning techniques that exploit features associated with reviews, such as ratings, helpful votes, and text content (4). A primary challenge in this approach is the lack of ground truth data with which to train and test models. In other words, in order to develop models that can identify fake reviews, one must first have a corpus of existing fake reviews (and real reviews) with which to train the model. Researchers have attempted to overcome the inherent challenge this poses largely by using laboratory-generated reviews and considering platform-filtered reviews as a proxy for fake reviews. Scholars have criticized these approaches as having serious limitations (4, 6–9). Lab-generated fake reviews may lack authenticity, and bot-generated or other low-quality reviews (those from non-English speakers or containing many grammatical errors), at best, pick the low-hanging fruit and miss the bulk of actual fake reviews. Moreover, using only reviews that are flagged and filtered by the platform's

Significance

Online reviews significantly impact consumers' decisions and are seen as crucial to the success of online markets. Despite this, the prevalence of fake reviews is arguably higher than ever, despite two decades of academic research on identifying and regulating them. We use data in which we directly observe which products buy fake reviews, and study how to identify them. We show that products buying fake reviews are highly clustered in the product reviewer network, due to their reliance on common reviewers. This allows us to detect them with high accuracy using both supervised and unsupervised methods. Unlike approaches relying on reviews' text, this approach is more robust to manipulation by sellers. Moreover, it is scalable and generalizable to many settings.

Author affiliations: [a]Anderson School of Management, University of California, Los Angeles, CA 90095; [b]Saunders College of Business, Rochester Institute of Technology, Rochester, NY 14623; and [c]Marshall School of Business, University of Southern California, Los Angeles, CA 90089

Author contributions: S.H., B.H., G.O., D.P., and A.T. designed research; S.H., B.H., G.O., D.P., and A.T. performed research; S.H., B.H., G.O., D.P., and A.T. analyzed data; and S.H., B.H., G.O., D.P., and A.T. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. A.G. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: brett.hollenbeck@anderson.ucla.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2211932119/-/DCSupplemental.

Published November 15, 2022.

algorithms, by definition, cannot progress our understanding of how to identify the large number of fake reviews that currently evade those filters (3), and methods trained on these data have been repeatedly shown to lack external validity (8).

In addition to the lack of high-quality ground truth data, another shortcoming of current methods is that sophisticated actors can potentially evade filtering algorithms trained on historical features. For example, methods relying on text analysis are fundamentally limited in that, even with sophisticated models, human reviewers have strong incentives to evade detection and will therefore strive to write fake reviews that are indistinguishable from organic reviews. If particular phrases or tendencies become used to filter reviews, they can then avoid using these in their reviews. More generally, extant approaches to detecting which products are manipulating their ratings have not grappled with the economic incentives of buyers and sellers. We use unique data in which we observe sellers buy fake reviews and demonstrate how these economic incentives can be harnessed to design an approach that can detect fake review buyers with high accuracy and with limited scope for evasion.

We argue that these problems can be overcome by focusing on the product reviewer network and identifying products that buy fake reviews rather than the fake reviews themselves. We hand-collect a large dataset that provides accurate ground truth data by directly identifying a large and representative set of products that buy fake reviews on Amazon.com. Using these data, we study the relative effectiveness of different approaches for detecting what products manipulate their ratings. We find that, compared to products not using fake reviews, products that use fake reviews are highly connected and clustered in the network, implying that products using fake reviews tend to have more reviewers in common than other products. This follows naturally from the fact that, while regular products receive their reviews from a dispersed set of millions of Amazon customers, products buying fake reviews must rely on the relatively small number of reviewers participating in the fake-review marketplace. Therefore, features derived from the product reviewer network are especially useful for regulating fake reviews because, in addition to being more predictive than text or metadata features, they are more difficult or costly to manipulate by sellers than text or review timing, because these features result directly from the inherent limitations of acquiring reviews from fake review marketplaces.

Data

To analyze fake review behavior on Amazon, we begin by collecting data from the private Facebook groups in which sellers buy reviews. While there are inherent limitations in observing fake-review activity, these groups are generally seen as the primary channel by which sellers find reviewers.

From March 2020 to October 2020, we identified about 23 fake review–related groups daily. These groups are large and quite active, each having about 16,000 members, on average, and 568 fake-review requests posted per day per group. Within these Facebook groups, sellers can obtain a five-star review that looks organic. Sellers post product pictures and review requests, after which the potential reviewer and the seller communicate via private Facebook messages. The vast majority of sellers that we observe buying fake reviews compensate the reviewer by refunding the cost of the product via a PayPal transaction after the five-star review has been posted, along with the cost of the PayPal fee, sales tax, and, in some cases, an additional commission. Reviewers are compensated for creating realistic seeming five-star reviews that evade Amazon's filters. This process differs from "incentivized reviews," where sellers offer free or discounted products or discounts on future products in exchange for reviews that disclose the transaction and are not required to be five stars (10).

Among these groups, and during our entire observation period, we observe Amazon sellers buying only positive reviews. As ref. 3 points out, this is likely because buying fake negative reviews to hurt competitors is costlier, as the buyer needs to incur the full cost of the competitor's product. In addition, the benefit of purchasing fake negative reviews is likely to be lower than that of buying a product's own positive reviews, because the shift in a product's own sales caused by a negative review on a competitor's page is indirect and dispersed across potentially many other competitor products besides the competitor for which the negative fake review is bought.

To identify which products are buying fake reviews, we hired a group of research assistants to visit these groups and select a random sample of products posted in them. We collect data from these random Facebook fake review groups using this procedure on a weekly basis from October 2019 to June 2020, and the result is a sample of roughly 1,500 unique products.

After identifying products whose ratings are manipulated, we collect data for these products on Amazon.com. We collect reviews and ratings for each of the products on a daily basis. For each review, we observe the rating, product ID, review text, review photos, and helpful votes.

In addition to collecting these data for the products buying fake reviews, we collect daily review data for a set of 2,714 competitor products to serve as a comparison set. To do so, for each product buying fake reviews, we select the two products that appear most frequently in the search results during a period covering 7 days before and after the date of the product's first Facebook post. The rationale is that we want to create a comparison set of products that are in the same subcategory as the products buying fake reviews and have a similar search rank before fake reviews are posted.

Identifying Which Products Buy Fake Reviews

There is an extensive literature on fake-review detection on online platforms such as Amazon, Yelp, and Tripadvisor. This literature has proposed methods for detection based on text features, image features, spatiotemporal differences, network features, sentiment, and others (see the many references in ref. 4). One of the main hurdles in creating algorithms capable of detecting fake reviews is the lack of ground truth data, that is, reviews that are known to be fake, and this literature has been criticized for relying on low-quality proxies (4).

We overcome this issue by focusing on the products that buy fake reviews. As described in *Data*, our dataset allows us to know, with a high degree of certainty, which Amazon products bought fake reviews, therefore providing accurate ground truth for our product-level analysis. We can then compare the performance of previously suggested methods and, in particular, test the performance of models that take advantage of the structure of the product reviewer network. Another recent study (11) has also utilized ground truth data obtained by monitoring Facebook groups to perform review- and product-level fake-review detection. However, the authors do not use the network structure as we do in this paper. Despite this, their qualitative analysis and interviews performed with buyers and sellers align with several of our findings, and the authors conclude that the network structure of the products buying fake reviews could be useful to detect products buying fake reviews.

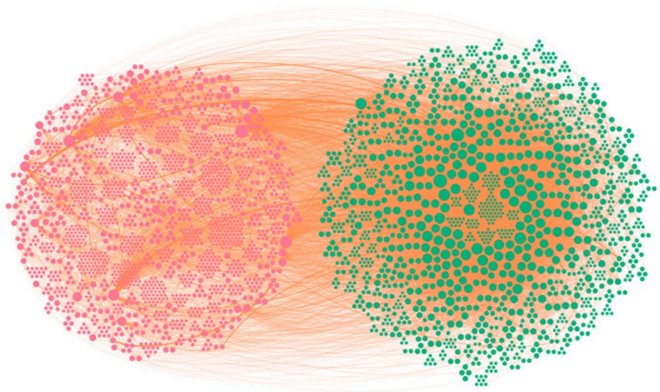

Fig. 1. The network structure of Amazon products in our dataset. Green-colored nodes are the fake review buyers. We filter the edges to show only the ones between pairs of products with $\geq$2 reviewers in common. A product's size is proportional to the number of reviewers in common with other products. The figure shows that fake review buyers are mainly the larger nodes in the network and are highly clustered, highlighting the dense connectivity between fake review buyers due to their participation in the fake-review marketplace.

Network Construction and Network Features Generation. Our approach to detecting sellers buying fake reviews exploits the network structure of Amazon products. We start by constructing a product network $G = (V, E)$ using our data, where G is the network, V is the set of nodes (i.e., products), and E is the set of edges. An edge between two products represents the existence of common reviewers. Fig. 1 shows the network structure of Amazon products in our dataset. Then, for each node (i.e., product), we compute its degree, eigenvector centrality, PageRank score, and clustering coefficient. We provide the mathematical details of how we compute these measures in *SI Appendix*, section 1.

The degree of a product is the total number of reviewers it has in common with other products in the network. Eigenvector centrality measures the structural importance by considering the product's proximity to other structurally important products in the network. A product's centrality score increases by having more reviewers in common with other products that themselves share many reviewers with other products in the network. PageRank, like eigenvector centrality, considers the importance of a product's neighbors when assigning a score to it. However, it mainly differs from eigenvector centrality by normalizing a neighbor product's structural importance by the number of its neighbors.

The above three measures (degree, eigenvector centrality, and PageRank) help us compare products' structural importance and understand how they relate to the overall network. In addition, we measure the products' connectivity within their neighborhood using the clustering coefficient. The clustering coefficient checks all pairs of a product's neighbors and considers how many of them have reviewers in common.

Additional Features. To compare the performance of different types of information at detecting fake-review buyers, we also generate features from available product metadata and review content such as review text, ratings, timestamps, and images.

The first set of features is generated from available metadata. These features include the number of reviews, average review rating, summary statistics of the time gap between reviews, the ratio of reviews that other consumers found helpful, the ratio of reviews with one-star review rating, the ratio of reviews with five-star review rating, the ratio of reviews that include review images, variation of review lengths for a product, and the average text similarity among the reviews of a product. These metadata should capture evidence of rating manipulation, such as having a disproportionate share of five-star reviews, excessive helpfulness votes, or odd timing characteristics such as long gaps in the arrival of reviews followed by the appearance of many at once.

We generate a second set of product-level features using review and product images. First, we use a pretrained deep convolutional neural network (12) for all images (both product and review) to extract the features as shown in *SI Appendix*, Fig. S1. We then calculate the angular similarity of all the images belonging to the same product, using the cosine distance between the vector representation of each image. Cosine distance is particularly effective for high-dimensional data (13). We calculate three sets of features that summarize the degree of image similarity between all product reviews (controlling for multiple images belonging to the same review), the similarities between the seller's product images and the review images, and the similarity among all pairs of review images. For each set of features, we calculate the minimum, maximum, average, and SD of the pair-wise similarities within a product.

The third set of product-level features generated from the review content is text features. We first combine the review bodies of a product and treat each product as a document. Then, we calculate the term frequency–inverse document frequency (TF–IDF) score of words in each document. We only consider the top 1,000 features that receive the highest TF–IDF scores for each product. Finally, each product is represented by a feature vector of length 1,000.

Table 1 shows all the product-level features we generate to consider various aspects of reviews, products, and reviewer behaviors to detect fake review buyers. *SI Appendix*, Table S1 shows the correlation coefficients between image, metadata, and network features.

Results

Supervised Approach. To test the predictive performance of the features we created, we employ a set of random forest classifiers.* Following the standard approach in the literature (14, 15), we train the classifiers on a random subset of 80% of the products and evaluate its prediction performance on the remaining 20%. We measure the area under the receiver operating characteristic curve (AUC), classification accuracy, true positive rate (TPR), true negative rate (TNR), and F1 score. The technical details of our estimated model and model building process are provided in *SI Appendix*, section 2.

Table 2 shows the performance of each type of feature set. We find that features constructed from the product reviewer network outperform metadata, text, and image features across all accuracy metrics. The combined model that includes all features performs only slightly better than the model using only network features. In addition to performing best on balanced measures of accuracy, the network feature model performs best on the true positive rate. This is especially important for a rating platform for whom avoiding false positives is an important goal.

Fig. 2 shows the relative importance of different individual features from the all-feature model in terms of their contribution to predictive power. The two most important features, by a large margin, are network features, specifically, the clustering coefficient and the eigenvector centrality score. The fact that these two features rank highest in importance highlights the reason why the network-based model performs so well. Products buying fake reviews appear to rely on a common set of reviewers, causing them to be more closely clustered in the product network compared to regular products.

*We also test other classifiers and obtain qualitatively similar results, which can be found in *SI Appendix*, section 3.

Table 1. Product-level features generated from review content and product network

Type	Feature	Description
Network	degree	Total number of reviewers in common with other products
	clustering coef	Clustering coefficient of the product
	eigenvector cent	Eigenvector centrality score of the product
	pagerank	PageRank score of the product
Metadata	tf-idf sim	Avg similarity of TF–IDF features between reviews of a product
	# reviews	Number of reviews
	avg review rating	Average review ratings
	time b/w reviews	Average, minimum, maximum, and SD of time in days between reviews
	share helpful	Share of reviews with helpful votes
	share 1star	Share of reviews with one-star rating
	share 5star	Share of reviews with five-star rating
	share photo	Share of reviews with review photo
	stdev review len	Variation of review lengths of a product
Image	img sim	Average, minimum, maximum, and SD of image similarity between product reviews
	sim review	Average, minimum, maximum, and SD of similarity between all pairs of review images
	sim product	Average, minimum, maximum, and SD of similarity between product and review images
Text	tf-idf	Top 1,000 TF–IDF features of a product

Table 2 shows that, when we test a more concise model using only these top two network features, it still outperforms all other models. This suggests that the predictive power of network features is sufficiently strong that, even when implemented in a simple way, they are more useful than models based on large sets of features derived from metadata, text, or images. We acknowledge that since platform techniques to detect fake reviews are not publicly disclosed, this method or something similar may already be used in practice.

Unsupervised Approach. While we have shown that a model trained on features derived from the product network is highly accurate at detecting fake review buyers, platforms may not have the ground truth data required to estimate this type of model. Therefore, to further validate the strength of a network-based approach, we extend our analysis to a much larger dataset of Amazon product reviews where we lack ground truth data and test whether network data allow us to identify, ex ante, which products are likely to be manipulating their ratings.

We use data that contain the entire universe of Amazon product reviews in the Home & Kitchen category, which contains about 65,000 products, 11 million reviews, and 6.1 million reviewers (16). Because we no longer know which products buy fake reviews in this larger dataset, we follow an unsupervised approach. We start by creating a product network as previously described. We then partition the products into 20 groups using the K-means clustering algorithm based on products' metadata and network features (17). Once the products are clustered, we test whether certain clusters are disproportionately likely to contain fake-review products. To test this, we use our pretrained classifier that was trained on all features, to estimate the proportion of products buying fake reviews in each cluster. Although this calculation would ideally be done using ground truth data, given the high out-of-sample accuracy of the all-features classifier, it should provide a good indication of the distribution of fake-review products among clusters.

Table 2. Out-of-sample prediction performance of the random forests classifier with varying sets of features

Features	AUC	Accuracy	TNR	TPR	F1 Score
Network	0.890	0.821	0.839	0.797	0.821
Top-two network	0.879	0.812	0.832	0.787	0.812
Metadata	0.874	0.811	0.858	0.748	0.810
Text	0.857	0.770	0.929	0.559	0.759
Image	0.592	0.599	0.792	0.343	0.577
All features	0.932	0.860	0.881	0.832	0.860

Fig. 3 shows the percentage and the total number of products identified as buying fake reviews in each cluster, using our classifier. We report summary statistics of clusters and the details of the unsupervised approach in *SI Appendix*, section 4.

This analysis shows that just two clusters (that contain only 3.4% of products) contain about 70% of the products identified as fake review buyers. In addition, within one of these clusters, a substantial majority of products are identified as buying fake reviews (83%). In *SI Appendix*, Table S4, we report the mean of the feature values in each cluster. The two clusters that contain the

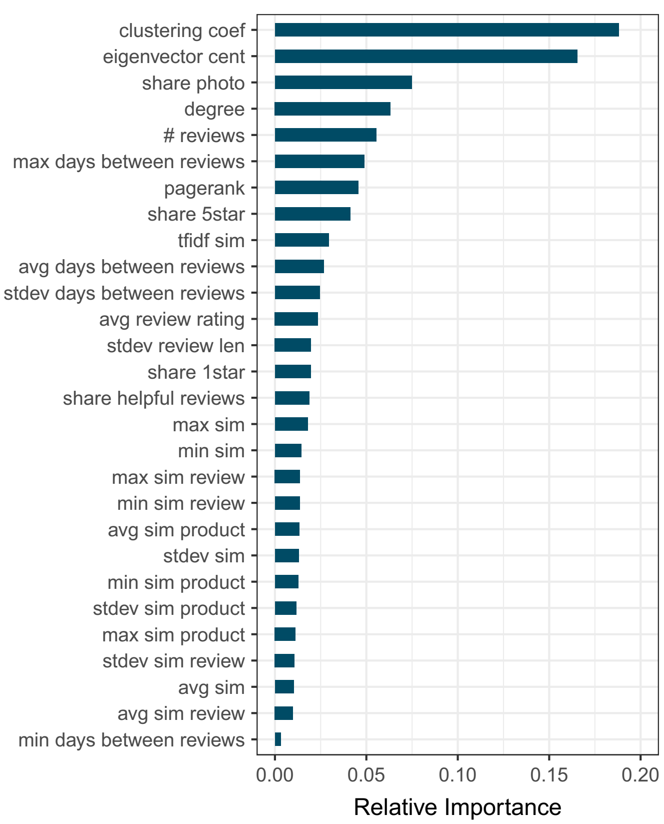

Fig. 2. Relative importance of features in terms of their contribution to prediction performance of the random forests classifier including all features.

 https://doi.org/10.1073/pnas.2211932119

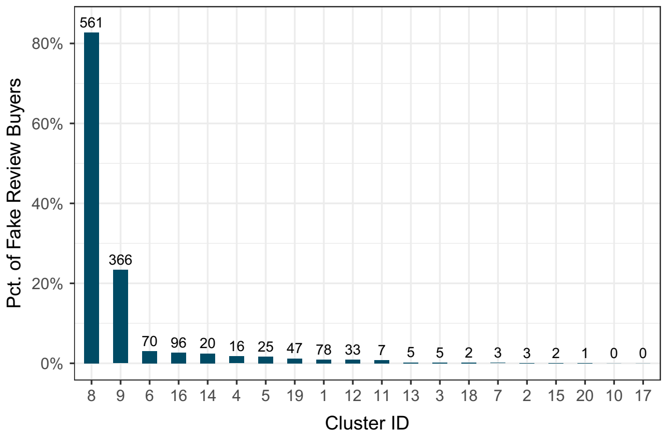

Fig. 3. The percentage of products that are identified as fake-review buyers by the random forest classifier in each cluster.

majority of fake-review buyers are highly distinctive, and only in terms of their network features. This suggests that a platform could use unsupervised methods to identify these tightly connected clusters of products without any ground truth data and use this information to identify likely fake review buyers.

Discussion and Conclusions

In this paper, we use a unique dataset and method to address a longstanding problem in the digital economy: rating manipulation. We use direct observation of what sellers buy fake reviews, and propose a network-based approach to detect these products.

Our analysis of the product network shows that products that buy fake reviews are highly clustered; while we don't have information about all Amazon reviews and reviewers, and we only see reviews that are ultimately posted, this seems to suggest that reviewers participating in the fake review marketplace are a relatively small part of the full universe of Amazon reviewers. Therefore, a classifier based on just two network features, clustering coefficient and eigenvector centrality, can identify products buying fake reviews with high accuracy, outperforming models trained on large numbers of features constructed from metadata, review text, and review images. In addition to being powerful features to detect products that buy fake reviews, a crucial implication of this insight is that network-based features are very costly to manipulate, because these features result directly from the inherent limitations of acquiring reviews from fake review marketplaces. Although the specific setting we study is Amazon.com, the disproportionate clustering of rating manipulators in the network structure is likely to hold across applications where fake reviews are found. This is because it will generally be the case that the set of reviewers used by rating manipulators is likely to be substantially smaller than the general population of reviewers, leading to the same pattern of unusual clustering observed in our network. Moreover, since we show that no ground truth data are necessary to identify these clusters, platforms can easily implement our approach. This insight may also contribute to the study of similar problems, such as botnet-directed ad fraud (18) and online review assessment (19).

Interestingly, our results complement existing methods used by platforms to mitigate rating manipulation. In particular, platforms are often more cautious with new users with few reviews than with experienced reviewers with many reviews. While this approach is probably correct in general and is supported by prior research (20), our results suggest that there exist a set of experienced reviewers posting many fake reviews, and that, by employing the network structure, the platform can detect them.

Finally, while we provide a method that can detect products that buy fake reviews, we do not address the question of what to do with this information. Amazon can target products, sellers, reviewers, or all of them. For example, it can remove all product reviews, remove the product from the platform, ban the seller from selling on the platform, or ban or sue fake reviewers. Although Amazon currently targets the writers of fake reviews, in our view, Amazon should target sellers. An approach would be to ban these sellers. However, in cases where the false-positive rate is nonzero, this can be quite costly. Furthermore, it could provide a perverse incentive for sellers to buy fake reviews for their competitors to try to get them banned. Therefore, a likely more balanced but still effective strategy would be to punish these sellers by reducing their visibility in the marketplace (e.g., by increasing their search rank) if they are suspected of buying fake reviews, which would offset and eliminate the financial incentive to do so.

More generally, our results show that the product network structure provides a robust source of information that can allow rating platforms to apply greater scrutiny to the products they list which, in turn, can deter sellers from committing fraud.

Data, Materials, and Software Availability. Data have been deposited in GitHub (https://github.com/dadepro) (21).

ACKNOWLEDGMENTS. We thank the Morrison Center for Marketing Analytics for generous funding. We thank seminar participants at the Johnson School of Management at Cornell University, the Ross School of Business at the University of Michigan, Pontificia Universidad Catolica de Chile, the Fox School of Business at Temple University, the Rochester Institute of Technology Marketing Workshop, the University of Southern California Marshall AI Workshop, and the 2022 Symposium on Statistical Challenges in Electronic Commerce Research for helpful comments.

1. S. Tadelis, Reputation and feedback systems in online platform markets. *Annu. Rev. Econ.* **8**, 321–340 (2016).
2. L. Molm, R. Hardin, M. Levi, Cooperation without trust? *Adm. Sci. Q.* **51**, 305–307 (2006).
3. S. He, B. Hollenbeck, D. Proserpio, The market for fake reviews. *Mark. Sci.* **41**, 896–921 (2022).
4. Y. Wu, E. W. Ngai, P. Wu, C. Wu, Fake online reviews: Literature review, synthesis, and directions for future research. *Decis. Support Syst.* **132**, 113280 (2020).
5. B. Dong, M. Li, K. Sivakumar, Online review characteristics and trust: A cross-country examination. *Decis. Sci.* **50**, 537–566 (2019).
6. A. Heydari, M. Tavakoli, N. Salim, Z. Heydari, Detection of review spam: A survey. *Expert Syst. Appl.* **42**, 3634–3642 (2015).
7. A. Mukherjee, V. Venkataraman, B. Liu, N. S. Glance, "What Yelp fake review filter might be doing?" in *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media* (Association for the Advancement of Artificial Intelligence, 2013), pp. 409–418.
8. M. Crawford, T. M. Khoshgoftaar, J. D. Prusa, A. N. Richter, H. A. Najada, Survey of review spam detection using machine learning techniques. *J. Big Data* **2**, 23 (2015).
9. D. U. Vidanagama, T. P. Silva, A. S. Karunananda, Deceptive consumer review detection: A survey. *Artif. Intell. Rev.* **53**, 1323–1352 (2019).
10. G. Burtch, Y. Hong, R. Bapna, V. Griskevicius, Stimulating online reviews by combining financial incentives and social norms. *Manage. Sci.* **64**, 2065–2082 (2018).
11. R. Oak, Z. Shafiq, The fault in the stars: Understanding the underground market of Amazon reviews. arXiv [Preprint] (2021). *arXiv: 2102.04217*. https://arxiv.org/abs/2102.04217v2 (Accessed 8 August 2022).
12. K. He, X. Zhang, S. Ren, J. Sun, "Deep residual learning for image recognition" in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, 2016), pp. 770–778.
13. S. L. France, J. D. Carroll, H. Xiong, Distance metrics for high dimensional nearest neighborhood recovery: Compression and normalization. *Inf. Sci.* **184**, 92–110 (2012).
14. M. Zhang, L. Luo, Can consumer-posted photos serve as a leading indicator of restaurant survival? Evidence from Yelp. *Manage. Sci.*, https://doi.org/10.1287/mnsc.2022.4359 (2022).
15. N. Kumar, D. Venugopal, L. Qiu, S. Kumar, Detecting review manipulation on online platforms with hierarchical supervised learning. *J. Manage. Inf. Syst.* **35**, 350–380 (2018).
16. J. Ni, J. Li, J. McAuley, "Justifying recommendations using distantly-labeled reviews and fine-grained aspects" in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, S. Padó, R. Huang, Eds. (Association for Computational Linguistics, 2019), pp. 188–197.
17. C. Fraley, A. E. Raftery, Model-based clustering, discriminant analysis, and density estimation. *J. Am. Stat. Assoc.* **97**, 611–631 (2006).
18. B. R. Gordon *et al.*, Inefficiencies in digital advertising markets. *J. Mark.* **85**, 7–25 (2021).
19. S. Luan, A. Mueen, M. Faloutsos, A. J. Minnich, "Online review assessment using multiple sources." US Patent 10089660 (2018).
20. M. Luca, G. Zervas, Fake it till you make it: Reputation, competition, and Yelp review fraud. *Manage. Sci.* **62**, 3412–3427 (2016).
21. D. Proserpio, Davide Proserpio, dadepro. GitHub. https://github.com/dadepro. Accessed 27 October 2022.