

Predictive Power for Product Sales: A Comparative Empirical Study of Consumer Reviews and Influencer Seeding Content

Yuqing Ye

*College of Economics, Hunan Agricultural University, Changsha, China
ye723012@gmail.com*

Abstract. Consumer reviews and influencer posts provide two different views of the same product. This study asks whether either source improves product-sales prediction beyond product attributes. The sample contains 90 skincare products, 3,060 reviews, and 450 influencer posts from Tmall and Taobao. Four OLS models compare a product-control benchmark with review, influencer, and combined feature sets. Category-stratified five-fold cross-validation evaluates held-out products, while HC3 robust standard errors are used for inference. The control model performs best (adjusted $R^2 = 0.651$, pooled out-of-fold $R^2 = 0.617$, RMSE = 0.470). Adding content features slightly raises unadjusted R^2 but reduces adjusted and out-of-sample performance, and no focal content coefficient is statistically significant. In this small, platform-concentrated sample, product attributes contain most of the available predictive information. Visible review and influencer signals do not improve sales prediction once controls are included. Because the sales window, temporal order, promotional exposure, and brand strength are unavailable, the findings describe predictive associations rather than causal effects.

Keywords: consumer reviews, influencer seeding, product sales, predictive modeling, category-stratified five-fold cross-validation

1. Introduction

Consumers leave two different digital traces around a product. One appears after purchase in ratings and reviews; the other is created by influencers through demonstrations, recommendations, and interaction with followers. Reviews may reduce uncertainty and provide social proof, whereas influencer content can attract attention and make product use more vivid. Because the two sources come from different actors and stages of the customer journey, each may capture a different aspect of demand. For firms allocating limited marketing resources, the useful question is not simply whether these signals correlate with sales, but whether they improve prediction after price, rating, category, and platform are considered.

Research suggests that both sources can matter, although their value varies. Meta-analyses link electronic word of mouth with market outcomes but find substantial differences across products, platforms, and measures [1, 2], and individual review text may add information beyond average ratings [3]. Influencer outcomes also depend on content, engagement, authenticity, relationships, and platform conditions [4, 5]. Livestream evidence links influencer characteristics to product sales [6],

but engagement is not the same as purchasing. Concentrated ratings, source credibility, campaign timing, and campaigns launched in response to weak sales may all reduce the extra information carried by these visible signals. These uncertainties make a common out-of-sample comparison useful.

Using the same 90 skincare products, we compare review-derived and influencer-derived features in four OLS models: a control benchmark, a review model, an influencer model, and a combined model. All models use the same controls and category-stratified validation folds. This allows the two sources to be compared under identical conditions and separates in-sample fit from prediction for held-out products. The analysis focuses on incremental predictive value rather than bivariate correlation and does not treat cross-sectional associations as causal effects.

2. Literature review and research framework

2.1. Consumer reviews and sales-related information

Online reviews carry ratings and textual detail. Ratings summarize evaluations, while comments describe attributes and experiences that a single score cannot capture. Reviews may reduce risk and provide social proof. Meta-analytic evidence links electronic word of mouth with sales, although effects vary by product, platform, and measurement [1, 2]. Individual reviews can affect purchase likelihood beyond aggregate ratings [3], while sentiment effects may depend on perceived credibility [7]. Because ratings cluster near the upper end in this sample, they may offer little discrimination. Reviews are therefore possible signals, not uniformly positive predictors.

Text analytics converts reviews into structured predictors through sentiment and topic information [8, 9]. Because the dataset fixes review volume at 34 observations per product, the analysis instead uses mean rating, rating dispersion, rating shares, dictionary sentiment, text length, and duplicate-text prevalence. These measures capture evaluation and textual characteristics, while aggregation may suppress individual variation. H1 tests whether this information set improves held-out prediction beyond product controls.

2.2. Influencer seeding content and consumer response

Influencers combine access to an audience with product endorsement and content creation, translating product attributes into platform-specific stories [4]. Their influence may come from reach, perceived authenticity, relationships with followers, and product demonstrations. A review of 214 studies finds that outcomes depend on the interaction of audience, influencer, brand, content, and technology characteristics [5]. Likes, comments, shares, and saves record attention and response, but they do not directly measure purchasing.

Influencer and content characteristics are associated with engagement, and the form of engagement varies with audience size [10]. Livestream research also links influencer characteristics to sales [6]. However, brands may increase seeding for weak-selling products, while entertaining posts may attract reactions without conversion. H2 therefore tests whether engagement, sentiment, text length, hashtag use, and promotional wording add predictive value after controls, without assigning a causal interpretation.

2.3. Research gap and hypotheses development

Studies usually examine reviews and influencer content separately. Sponsored, marketer-generated, and user-generated content produce different brand responses [11], and user-generated brand content can increase purchase intention [12]. These findings do not show which source better predicts sales. Few studies compare both feature sets with common controls and the same held-out products. In a small sample, in-sample R^2 may favor larger models that fit noise. We therefore compare adjusted fit and identical out-of-sample validation.

Product attributes form the benchmark. The hypotheses test whether review features improve held-out prediction, whether influencer features improve it, and whether the two sources differ in predictive performance:

H1: Consumer-review features provide incremental out-of-sample predictive value beyond product controls.

H2: Influencer-seeding features provide incremental out-of-sample predictive value beyond product controls.

H3: The out-of-sample predictive performance of consumer-review features differs from that of influencer-seeding features.

3. Research methodology

3.1. Data and sample construction

The files contain 90 skincare products, 3,060 reviews, and 450 influencer posts. Each product has 34 sampled reviews and five posts; each of five categories has 18 products. The sample includes 81 Tmall and nine Taobao products, limiting generalization. Fixed content counts indicate capped sampling, but the selection rule is undocumented. Sales range from 1,500 to 50,000 units ($M = 11,814.44$, $SD = 10,496.54$), and prices from CNY 29 to CNY 2,680. Because the sales window and temporal order are unknown, results are interpreted only as product-level predictive associations.

Identifiers, numeric fields, missingness, integrity, and IQR outliers were checked; no records required deletion. Duplicate review text affected 103 records (3.37%), but removing copies barely changed product means, so all records were retained. Sales and skewed engagement variables were transformed as $\ln(1 + x)$.

3.2. Feature extraction

Review features covered ratings, rating shares, sentiment, text length, and duplicate-text share. Influencer features covered engagement, text length, sentiment, hashtag count, and promotional-word count. Sentiment equals matches to 15 positive terms minus matches to 14 negative terms; it is a transparent but coarse dictionary measure, not a validated Chinese sentiment model.

3.3. Statistical analysis and predictive evaluation

Descriptive statistics characterize both sources, and Spearman correlations assess skewed bivariate associations. Four OLS models predict log sales: M0 uses log price, product rating, category, and platform; M1 adds five review indicators; M2 adds five influencer indicators; M3 combines both sets. R^2 and adjusted R^2 assess in-sample fit. Category-stratified five-fold cross-validation (seed 42) evaluates identical held-out products by pooled out-of-fold R^2 , RMSE, and MAE. A paired sign-flip test (200,000 draws; seed 42) compares M1 and M2 absolute out-of-fold errors. HC3 robust

standard errors support inference. Coefficients from the 90-product cross-section are conditional associations, not causal effects.

4. Results and discussion

4.1. Descriptive patterns and source differences

Table 1. Descriptive statistics

Variable	N	Mean	Std. Dev.	Median	Range
Sales volume (units)	90	11,814.44	10,496.54	8,000.00	1,500–50,000
Price (CNY)	90	269.03	366.05	159.00	29–2,680
Product rating (1–5)	90	4.73	0.10	4.70	4.60–4.90
Review rating (1–5)	90 products	4.34	0.18	4.35	3.68–4.74
Review text length	90 products	33.99	1.35	34.09	30.94–37.59
Influencer text length	90 products	48.63	2.14	48.70	43.40–55.80
Mean influencer likes per post	90 products	24,588.91	6,839.24	24,552.70	6,507–39,702
Total influencer engagement	90 products	169,214.59	48,087.86	168,986.00	44,885–283,479

Notes. Consumer review metrics are aggregated at the product level based on 34 sampled reviews per product.

Table 1 shows wide variation in sales and price but little variation in ratings: 58.17% of reviews are five-star and 84.54% are four- or five-star. Influencer posts are 14.64 characters longer than reviews ($t = 70.33$, $p < 0.001$), while their sentiment difference is not significant ($t = 1.46$, $p = 0.147$).

4.2. Bivariate associations with sales




Figure 1. Consumer-review indicators and log sales

Note. Lines are descriptive fitted trends with confidence bands; they do not imply causation.

Figure 1 shows no significant monotonic relationship between sales and mean review rating (Spearman $\rho = -0.114$, $p = 0.285$) or mean sentiment ($\rho = 0.030$, $p = 0.781$). This result does not

imply that reviews are irrelevant to shoppers. Instead, aggregation to 90 product means and the clustering of ratings near the upper end may leave too little variation to distinguish sales levels.

Total influencer engagement is negatively associated with sales in the bivariate analysis (Spearman $\rho = -0.291$, $p = 0.005$), but the data do not reveal why. Brands may support weak-selling products, campaigns may follow changes in sales, or engaging content may attract attention without converting it into purchases. The result therefore does not show that influencer activity reduces sales. Price has a stronger negative association with sales ($\rho = -0.461$, $p < 0.001$).

4.3. Comparative predictive performance

Table 2. OLS model performance: in-sample fit and category-stratified five-fold prediction

Metric	M0: Controls	M1: + Reviews	M2: + Influencer	M3: + Both
In-sample R^2	0.678	0.686	0.692	0.702
Adjusted R^2	0.651	0.637	0.643	0.632
Mean fold R^2 (SD)	0.588 (0.079)	0.523 (0.158)	0.525 (0.094)	0.469 (0.159)
Pooled OOF R^2	0.617	0.575	0.562	0.526
Pooled OOF RMSE	0.470	0.494	0.502	0.523
Pooled OOF MAE	0.390	0.409	0.417	0.437

Notes. The dependent variable is $\log(1 + \text{sales volume})$. Higher R^2 indicates better model fit, lower RMSE and MAE indicate higher prediction accuracy. All models include category and platform dummy variables.

Table 3. OLS regression coefficients with HC3 robust p-values

Variable	M0: Product Control	M1: + Consumer Reviews	M2: + Influencer Seeding	M3: + Both Sources
Log price	-1.131*** (<0.001)	-1.147*** (<0.001)	-1.075*** (<0.001)	-1.088*** (<0.001)
Product rating	6.242*** (<0.001)	6.397*** (<0.001)	5.790*** (<0.001)	5.940*** (<0.001)
Mean review rating	—	0.570 (0.331)	—	0.567 (0.344)
Review sentiment	—	-0.155 (0.551)	—	-0.211 (0.406)
Review text length	—	-0.027 (0.575)	—	-0.034 (0.509)
Log influencer engagement	—	—	-0.237 (0.214)	-0.241 (0.232)
Influencer sentiment	—	—	0.050 (0.651)	0.087 (0.450)
Influencer text length	—	—	-0.004 (0.904)	-0.001 (0.971)

*Notes. HC3 robust p-values are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Category and platform dummy variables are included in all models but omitted for brevity. Review-rating dispersion and low-rating share are included in M1 and M3, while influencer-engagement dispersion and promotional-word count are included in M2 and M3; these four coefficients are also omitted for brevity. The dependent variable is $\log(1 + \text{sales volume})$.*

The simplest specification performs best in Table 2. M0 has an in-sample R^2 of 0.678, an adjusted R^2 of 0.651, and a pooled out-of-fold R^2 of 0.617 (RMSE = 0.470; MAE = 0.390). Table 3 indicates a negative association between log price and log sales ($b = -1.131$, $p < 0.001$) and a positive association between product rating and log sales ($b = 6.242$, $p < 0.001$). These estimates are not causal because the data do not include brand strength, product age, promotion, or exposure.

H1 is not supported. Adding review variables raises unadjusted R^2 from 0.678 to 0.686, but adjusted R^2 falls to 0.637 and out-of-fold R^2 to 0.575. All review variables have HC3 p-values above 0.33 in M1. Concentrated positive ratings, fixed sampling of 34 reviews per product, and product-level aggregation may leave little incremental signal.

H2 is not supported. M2 remains below M0 in adjusted R^2 (0.643) and out-of-fold R^2 (0.562); total influencer engagement is non-significant after controls ($b = -0.237$, $p = 0.214$). M3 has the highest unadjusted R^2 (0.702), the lowest adjusted R^2 (0.632), and the weakest out-of-fold performance ($R^2 = 0.526$, RMSE = 0.523). H3 is also not supported: M1 slightly exceeds M2 in out-of-fold R^2 (0.575 vs. 0.562) and RMSE (0.494 vs. 0.502), but their paired absolute errors do not differ significantly (sign-flip $p = 0.655$).

5. Conclusion

For the 90 skincare products examined here, product controls provide more predictive information than either consumer reviews or influencer seeding. M0 achieves an adjusted R^2 of 0.651 and a pooled out-of-fold R^2 of 0.617; adding either or both content sources does not improve prediction for held-out products.

Price is negatively associated with sales and product rating positively associated, but no focal review or influencer indicator is significant after controls. Influencer visibility should not be treated as a reliable demand proxy, nor favorable review averages as automatic incremental information about sales.

The findings remain exploratory. The sample is small and Tmall-dominant, uses fixed content counts and dictionary sentiment, and lacks the sales window, temporal order, brand strength, promotion, exposure, creator reach, and product age. These limits preclude causal claims and broad generalization. Future studies should use natural content volumes, broader platforms, validated Chinese-language text measures, richer controls, and content observed before a defined sales period. Repeated or nested validation with uncertainty intervals could assess the stability of small predictive differences.

References

- [1] Babić Rosario, A., Sotgiu, F., De Valck, K., & Bijmolt, T. H. A. (2016). The effect of electronic word of mouth on sales: A meta-analytic review of platform, product, and metric factors. *Journal of Marketing Research*, 53(3), 297–318.
- [2] You, Y., Vadakkepatt, G. G., & Joshi, A. M. (2015). A meta-analysis of electronic word-of-mouth elasticity. *Journal of Marketing*, 79(2), 19–39.
- [3] Vana, P., & Lambrecht, A. (2021). The effect of individual online reviews on purchase likelihood. *Marketing Science*, 40(4), 708–730.
- [4] Campbell, C., & Farrell, J. R. (2020). More than meets the eye: The functional components underlying influencer marketing. *Business Horizons*, 63(4), 469–479.

- [5] Joshi, Y., Lim, W. M., Jagani, K., & Kumar, S. (2025). Social media influencer marketing: Foundations, trends, and ways forward. *Electronic Commerce Research*, 25, 1199–1253.
- [6] Zhang, S. K., Tang, T. Y., & Krallman, A. (2024). Navigating livestream commerce: A dual-lens framework of influencers' impact on product sales. *Journal of Business Research*, 185, 114919.
- [7] Wang, Q., Zhang, W., Li, J., Mai, F., & Ma, Z. (2022). Effect of online review sentiment on product sales: The moderating role of review credibility perception. *Computers in Human Behavior*, 133, 107272.
- [8] Yang, N., Korfiatis, N., Zissis, D., & Spanaki, K. (2024). Incorporating topic membership in review rating prediction from unstructured data: A gradient boosting approach. *Annals of Operations Research*, 339, 631–662.
- [9] Li, X., Wu, C., & Mai, F. (2019). The effect of online reviews on product sales: A joint sentiment-topic analysis. *Information & Management*, 56(2), 172–184.
- [10] van der Harst, J. P., & Angelopoulos, S. (2024). Less is more: Engagement with the content of social media influencers. *Journal of Business Research*, 181, 114746.
- [11] Müller, J., & Christandl, F. (2019). Content is king—But who is the king of kings? The effect of content marketing, sponsored content & user-generated content on brand responses. *Computers in Human Behavior*, 96, 46–55.
- [12] Mayrhofer, M., Matthes, J., Einwiller, S., & Naderer, B. (2020). User generated content presenting brands on social media increases young adults' purchase intention. *International Journal of Advertising*, 39(1), 166–186.